

# Ray Solomonoff and the Dartmouth Summer Research Project in Artificial Intelligence, 1956

Grace Solomonoff\*

Oxbridge Research

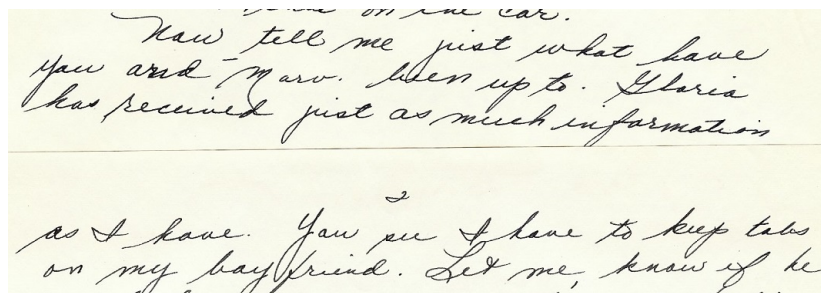
Mailing Address: P.O.B. 400404, Cambridge, Ma. 02140, U.S.A.  
prettyvivo@gmail.com <http://raysolomonoff.com>

## Abstract

This is the story of Ray Solomonoff and the Dartmouth Summer Workshop in 1956, often called the beginning of Artificial Intelligence. His notes and those of his friends show the events that led to it, what that amazing summer was like, and what happened after.

## Introduction: The Mysterious Science

A letter to Ray Solomonoff from his girlfriend Louise, sent in July 1956, asks about a mystery:



Now tell me just what have you and Marv been up to. Gloria has received just as much information as I have. You see I have to keep tabs on my boy friend. Let me know if he

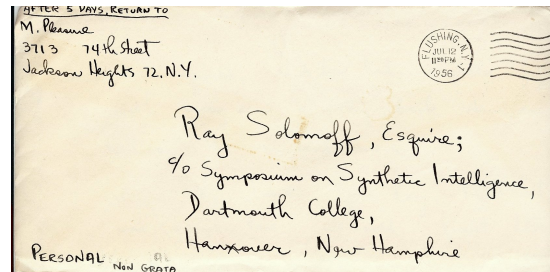
“Now tell me, just what have you and Marv been up to — Gloria has received just as much information as I have...”

What indeed? Gloria is the wife of the famous mathematician Marvin Minsky, then a Harvard Junior Fellow in Math and Neurology. Ray Solomonoff was a graduate of the University of Chicago, then working at Technical Research

---

\*Thanks to Alex Solomonoff, Donald Rosenthal, Sally Williamson, Julie Sussman for their invaluable help. Thanks to Professors Carey Heckman and James Moor for sharing copies of the Rockefeller Foundation file on the Dartmouth conference; to Trenchard More and Marvin Minsky for sharing files and interviews; to John McCarthy for interviews.

Group in New York. Marvin and Ray were spending the summer at Dartmouth College with a group of scientists, and what they were up to was research on an unusual new science — People didn't agree on what it was, how to do it or even what to call it:



Envelope of another letter sent to Ray during the Dartmouth Summer.

Some other names were cybernetics, automata theory, complex information processing, [16, p. 115] or just “thinking machines.”

But whatever the name of this strange new discipline, a group of scientists had gathered at the Dartmouth campus in a brainstorming session: how to develop a Machine that Thinks.

Since that time, some see the Dartmouth Summer Research Project of 1956 as a milestone, the event that initiated Artificial Intelligence (AI) as a research discipline. [44, p. 149][19, p. 87] To others it is just “a piece of AI mythology.” [8, p. 5] It did not achieve the collective results the organizers aimed for. But whatever it was, it made a lasting effect on AI.

My report describes what occurred during that summer, and some ways it impacted AI. Ideas were shared that influenced each other, in ways not commonly known. The focus is on one scientist who was part of the AI community, Ray Solomonoff. As much as possible, my sources are documents and reminiscences by members of the group and in particular Ray's handwritten notes. The first part looks more at events; the second, more at ideas.

Part One briefly describes Ray's early life before the Dartmouth Project. Then it looks at how the project was organized and at daily life during the eight weeks of the Summer Workshop. Some detailed quotes show how diverse the discussions were and what problems challenged the group. The notes by Ray and another participant, Trenchard More, provide most of the observations about daily life during the summer.

Part Two: Though the Summer Workshop never achieved its vision of a thinking machine, it did contribute ideas. This report looks at three areas: a shift in AI orientation from semantic, neural net and numeric processes, to symbolic processes; the growth of systems that solve highly specific problems rather than ones that focus on general learning; and a nascent division between deductive logic-based systems and inductive probabilistic-based systems.

I look in detail at two specific ideas that are oriented toward Ray. First, Ray's and Marvin's notes show that Ray influenced the change toward symbolic

systems in a way not well known by the AI community. Second, Ray's own ideas about probability were influenced by John McCarthy.

Ray, Marvin, and McCarthy were the only three who spent the whole time at the Workshop, and I have more notes about them than the other participants except for Trenchard More, who also had many notes. My main source is notes by Ray, so this particular report has a "Ray's eye view" of that summer.

But all of the participants made wonderful contributions to the scientific world. The summer reaffirmed something: uncensored thought shared in a free environment inspires the creative process. Recently uncovered notes and recorded conversations by participants at the Dartmouth Summer Workshop will show how.

## Part I: Intelligence as Never Before

### Ray and thinking machines — the early years

#### Ray's Early Life

Ray was the son of Russian Jewish immigrants, Julius Solomonoff and Sarah Mashman, who had escaped from the unrest and discrimination of 1910–1917 during the fall of Czar Nicholas before the Bolshevik revolution. They met and married in New York where their first son, George, was born in 1922. After they moved to Cleveland, Ohio, their second son, Ray, was born on July 25, 1926.



Left to Right: Ray and his father. Ray (on the left), with his brother George.  
Ray on a horsie.

Ray loved math and science from his earliest years, which his parents encouraged. He graduated in 1944 from the progressive Glenville High School,

and after a two-year stint in the Navy, entered the University of Chicago.


Those were brilliant years at the University of Chicago. Enrico Fermi's students taught Ray from their compendium of Fermi's notes (some titled "Unclear Notes" rather than "Nuclear Notes"). Anatole Rapoport was his teacher and Rudolf Carnap, who was trying to find a way to describe the universe by a digital string, was his professor of Philosophy.[2]

A fellow Navy man, Kirby Smith, was Ray's roommate in 1950. Kirby wrote in a letter:

"When I first met Ray he was sitting in an armchair surrounded by piles of manuscript covered in symbolic logic, to which he was continuing to add at a furious rate. It was some time before I took in the significance of all that." [7]


In a letter written in 1950, Ray tells a girlfriend that cybernetics (a common name for thinking machines) has been his chief interest for the past 4-5 years.

Here's what Ray was reading:



## The Tenor of the Time

Science fiction with its expectations for the future had blossomed since Hugo Gernsback's publication in 1926 of "Amazing Stories" and John Campbell's "Astounding" magazine, begun in 1930. Sci Fi fandom had recently been born, and Ray was president of the newly formed, always cash-strapped, Sci Fi club at the University. Thinking machines were featured in Sci Fi:



Science fiction from 1953, 1954 and 1957.

In 1941 the electronic computer had been introduced to the public. By the 1950s, scientists were debating how to use computers for thinking machines.

Both government and private companies funded new technology. After World War II money shifted from direct war needs to consumer and other products. Nuclear power came on in a big way to outperform British, French and Italian efforts. With the rise of the cold war, competition with Russia led to more space research,[48] while the Department of Defense gave grants for broadly defined research in computers and robotics, funding many civilian projects.[26, p. 203]

These were some of the elements that came together, making the time right for thinking machines.

After graduation Ray went to work in New York. At a conference, he met Marvin Minsky and they became lifelong friends.



Ray, a friend, and Marvin at Ray's apartment, in the early 1960s.

Marvin was one of the four who organized the Dartmouth conference. Due to Marvin, and to his own interests, Ray was one of the 11 original planned attendees to the Dartmouth Summer Workshop.

## Planning the summer

How did this unusual summer project come into being?

### An Inspired Decision

In the 1950s, those early days of computing and robotics, there was much confusion and some rancor about what thinking machines and robots would be like. Some, like Norbert Wiener, believed they would be humanoid, some did brain modeling, some focused on semantic-based systems, others on mathematical logic.

Among the confusion and politics, early in 1955 John McCarthy, a young Assistant Professor of Mathematics at Dartmouth, recognized the incredible



potential of this — whatever it was — and the possibilities for some kind of coherence about it.

He boldly picked the name “Artificial Intelligence.” He chose the name partly for its neutrality; avoiding a focus on narrow automata theory, and avoiding cybernetics which was too heavily focused on analog feedback, as well as him potentially having to accept the assertive Norbert Wiener as guru or having to argue with him.[26, p. 53] He began pushing to get a group together. McCarthy hoped a special conference would clarify direction and forge ahead.

### **The Proposal**

In February, 1955, McCarthy first approached the Rockefeller Foundation to request funding for a summer seminar at Dartmouth for about 10 participants. In June, he and Claude Shannon, a founder of Information Theory then at Bell Labs, met with Robert Morison, Director of Biological and Medical Research at the foundation, and discussed what kind of proposal might get funded.

At the meeting with Morison, participants were suggested such as Nat Rochester, head of Information Research at I.B.M., a designer of the early IBM 701 computer, Marvin Minsky, John Von Neuman, and others. Morison, however was very unsure that money would be made available for such a visionary project.[24]

McCarthy pressed ahead to prepare a proposal for the summer. He was enthusiastically joined by Shannon, Marvin Minsky, and Nat Rochester. On September 2, 1955, the four sent a proposal to the Rockefeller Foundation:

“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it . . . We think that a significant advance can be made . . . [on components of intelligence] . . . if a carefully selected group of scientists work on it together for a summer. . .”[15, p. 2] and “For the present purpose the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving.”[26, p. 53]<sup>1</sup>

Topics to study: automatic computers and programs for them; programming computers to use a language; neuron nets; machine self-improvement; classifying abstractions; and two areas particularly relevant to Solomonoff’s work, though they would not be part of AI for many years — how to measure the complexity of calculations, and randomness and creativity.[15, pp. 2–5]

### **Gathering the Group together**

At the meeting with Morison, Shannon and McCarthy had suggested other participants: Donald Hebb, Professor of Psychology at McGill, and from Britain, Ross Ashby of “Design for a Brain” fame and possibly Donald MacKay, a founder

---

<sup>1</sup>not in Ray’s copy.

of the British school of information theory. Oliver Selfridge who was connected with Norbert Wiener's cybernetics group was also proposed,[24] Minsky knew Oliver Selfridge well, and strongly supported his inclusion. McCarthy met with Herb Simon and Allen Newell at Carnegie Tech and discussed their work developing a program to find proofs for theorems in Principia Mathematica. They both agreed to attend [11]

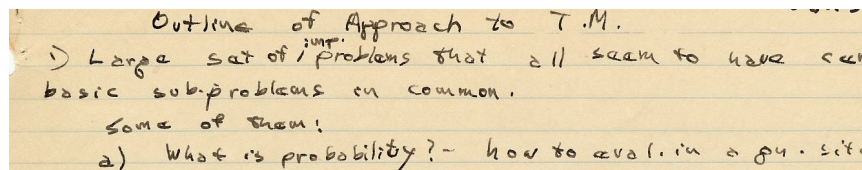
Among the scientists and mathematicians, Minsky suggested Ray Solomonoff calling him, a "best buy." He wrote the other organizers: "It has come to my attention that R. Solomonoff incessantly fills notebooks with resumes and observations on his and other peoples ideas, including those with which he may disagree. I cannot begin to describe the extraordinary depth and completeness of these notes. On this ground alone he would be a great asset when it comes to writing a report of the project; if Ray has been present, the bulk of the work may be done automatically." [18]

John McCarthy hoped the summer would produce some specific result. The group of scientists was motley, however, and so were their proposals. They included Shannon's idea for application of information theory to computing machines; Minsky's plan relating successful machine output activity to its ability of assembling "motor abstractions" as the environment changes; Rochester's program ideas for originality which included use of randomness; More's work on developing a programming technique for handling axiomatic systems in algebraic form.[15, pp. 2-5]

Everyone had a different idea, a hearty ego, and much enthusiasm for their own plan. Forging ahead together on anything at all would require some fancy footwork! But everyone was hopeful: Marvin wrote

"... by the time the project starts, the whole bunch of us will, I bet, have an unprecedented agreement on philosophical and language matters so that there will be little time wasted on such trivialities." [18]

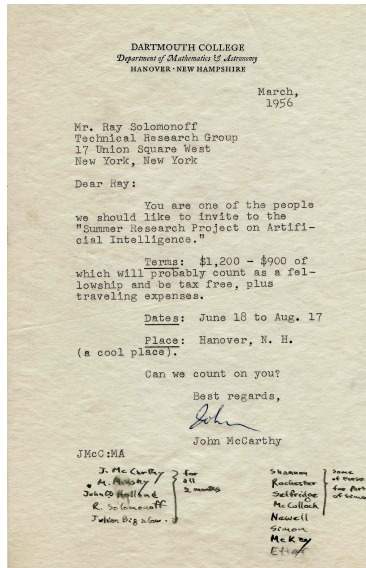
Here is the beginning of Ray's outline for work in thinking machines at the Dartmouth Summer, sent in March 1956:



What is probability?

Probability and its relation to prediction became Ray's main focus for the rest of his life. But at that time, it wasn't known how probability could be useful in AI.[41, p. 1]

Ray was invited,



On his invitation he noted some of the others who planned to attend.

On May 26, 1956, McCarthy notified Robert Morison of the 11 attendees:  
**for the full period:** 1) Dr. Marvin Minsky 2) Dr. Julian Bigelow 3) Professor D.M. Mackay 4) Mr. Ray Solomonoff 5) Mr. John Holland 6) Mr. John McCarthy.

**for four weeks:** 7) Dr. Claude Shannon, 8) Mr. Nathaniel Rochester, 9) Mr. Oliver Selfridge.

**for the first two weeks:** 10) Mr. Allen Newell and 11) Professor Herbert Simon.

He noted, “We will concentrate on a problem of devising a way of programming a calculator to form concepts and to form generalizations. This of course is subject to change when the group gets together.”[10]

The actual participants came at different times, mostly for much shorter times. Trenchard More replaced Rochester for three weeks, and MacKay and Holland did not attend — but the project was set to begin.

So around June 18, 1956, the earliest participants (perhaps only Ray, maybe with Tom Etter) arrived at the Dartmouth campus in Hanover, N.H., to join John McCarthy who already had an apartment there. Ray and Marvin stayed at Professors’ apartments, but most would stay at the Hanover Inn.





## The Summer of 1956

### When Did It Happen?

The Dartmouth Workshop is often said to have run for six weeks in the summer of 1956.[26, p.53] Ray's notes, however, say the workshop ran for "roughly eight weeks, from about June 18 to August 17." [34] Marvin, and Trenchard's notes as well, agree.[22] Ray's Dartmouth notes start on June 22; June 28 mentions Minsky, June 30 mentions Hanover, N.H., July 1 mentions Tom Etter. On August 17, Ray gave a final talk, so the event ran about eight weeks.

### Who Was There?

It has been thought that nobody had an accurate list of who actually was there during that time. So many people came for such short times, or had to cancel their planned attendance that there has been considerable confusion about this. Fortunately Ray's notes state who actually came as well!

#### Lists of Attendees:

**McCarthy's List:** Unfortunately McCarthy lost his list!<sup>2</sup>

**Trenchard's List:** Trenchard wrote, "Arranged roughly in order of the person's degree of attendance, interest, centrality or suggested attendance at the Project as perceived from my being in residence during some of the weeks."

- |               |               |                      |                         |
|---------------|---------------|----------------------|-------------------------|
| 1) McCarthy   | 9) Newell     | 17) Sayre            | 25) MacKay              |
| 2) Solomonoff | 10) Simon     | 18) Shoulders        | 26) Uttley              |
| 3) Minsky     | 11) Milner    | 19) Woodrow(sic)     | 27) Frankel             |
| 4) Rochester  | 12) McCulloch | 20) Nash             | 28) Holland, John       |
| 5) Shannon    | 13) Ashby     | 21) Bernstein        | 29) Pitts, Walter       |
| 6) Selfridge  | 14) Robinson  | 22) Backus           | 30) Wiener, Norbert     |
| 7) More       | 15) Etter     | 23) Moore E.F.       | 31) Rappert             |
| 8) Bigelow    | 16) Samuel    | 24) Hagdbarger, D.W. | 32) Shipero, Norman[22] |

In a later letter by Herbert Simon, he's quoted as also mentioning Professor F.B. Fitch.

<sup>2</sup>Instead, after the Dartmouth Project McCarthy sent Ray a preliminary list of participants and visitors plus those interested in the subject. There are 47 people listed.[6]

**Ray's List:** is smaller, but accurate for actual attendance:

- |                   |                     |                      |                            |
|-------------------|---------------------|----------------------|----------------------------|
| 1) Solomonoff     | 6) Nat Rochester    | 11) Abraham Robinson | 16) Shoulders              |
| 2) Marvin Minsky  | 7) Oliver Selfridge | 12) Tom Etter        | 17) Shoulder's friend      |
| 3) John McCarthy  | 8) Julian Bigelow   | 13) John Nash        | 18) Alex Bernstein         |
| 4) Claude Shannon | 9) W. Ross Ashby    | 14) David Sayre      | 19) Herbert Simon          |
| 5) Trench More    | 10) W.S. McCulloch  | 15) Samuel           | 20) Allen Newell[39, p. 1] |

#### **A Few More Attendees:**

It has been said that Bigelow and Shannon didn't attend, but Ray's notes show that both did, and that Bigelow spoke on August 15. Ray doesn't mention Bernard Widrow, but apparently he visited, along with W.A. Clark and B.G. Farley.[8, p. 5] Trenchard mentions R. Culver and Ray mentions Bill Shutz. Herb Gelernter didn't attend, but was influenced later by what Rochester learned.[26] Gloria Minsky also commuted there (with their part-beagle dog, Senje, who would start out in the car back seat and end up curled around her like a scarf), and attended some sessions (without Senje).

#### **Daily Attendance**

Ray, Marvin and John McCarthy were the only three who stayed for the full eight weeks. Trenchard later said, "I attended three of the weeks, which is probably more than anyone else there, except Ray, who was there most — nearly all of the time, and I think Ray was there even more often than Minsky." Trenchard took attendance during two weeks of his visit.[23] From three to about eight people would attend the daily sessions.

#### **Daily Attendance Examples:**

**The whole fifth week, 16-20 July:** Etter, McCarthy, More, Solomonoff;

**6th week, 23-27 July: Monday:** Ashby, McCarthy, Minsky, More, Solomonoff

**Tuesday:** Ashby, Culver, McCarthy, McCulloch, Minsky, More, Solomonoff

**Wednesday:** McCarthy, Minsky, More, Solomonoff

**Thursday:** McCarthy, Minsky, More, Robinson, Solomonoff

**Friday:** McCarthy, More, Robinson.[21, p. 1]

#### **Where Did They Meet?**

They had the entire top floor of the Dartmouth Math Department to themselves, and most weekdays they would meet at the main math classroom where someone might lead a discussion focusing on his ideas.

They tried to find meanings in common. One afternoon they looked up the word "heuristic" in the dictionary . . .

. . . In 2006, a Conference, "AI@50," was being planned at Dartmouth to honor the summer of 1956. No one knew exactly where the workshop group met. Trenchard met with the organizers, Professors James Moor and Carey Heckman, and guided them to the building:

Trenchard said "...they photographed me walking up the stairs and I explained where the rooms were where McCarthy was and said there was a dictionary in the room where we looked up the word heuristic and Wendy looked around and said 'Oh there's a dictionary in a stand over there.' I said 'that's the dictionary.' And so that's how even though the math department had moved out and another department had moved in, the dictionary was still there and on the stand. It was a big fat dictionary." [23]



Window looking out from the math room, and the dictionary, still open in 2014 at the word "heuristic."

### What Did They Say?

The participants shared widely different views and many described their own projects, usually relating them to ways that might make a machine able to learn, solve problems, become more intelligent. The atmosphere was more of discussion than talks. Different attendees would lead the discussion, including Ashby, Robinson, More, Samuel, McCarthy, Bernstein, Minsky, Bigelow, as well as Ray.

**An Example: Ashby's Talk:** Both Trenchard and Ray took notes of Ashby's talk, July 23, 1956. The following is a fragment:

**Trenchard's notes:** Ashby mentions his new book 'Introduction to Cybernetics;' ... Ashby stresses: "a simple machine appears to be extraordinary when viewed psychologically. When part of a mechanism is concealed from observation, the behavior of the machine seems remarkable."

**Both:** Ashby describes homeostats [electromechanical machines that could adapt to a changing environment].

**Trenchard's notes:** In the original homeostat, if the system was stable, the [magnetic] needles of the four units of Ashby's machine would stay in the center. If unstable, one or more of the needles would diverge to hit a stop. When this happened, the machine

would change one or more of the uniselectors, reorganizing the input connections, and the homeostat would search the new field for stability.


**Ray's notes:** Problem 1). No memory of previous solutions. i.e. learns to handle A, then learns B -> then going back to A takes as much time as before.

**Trenchard's notes:** Problem 2). time for problem-solving goes up exponentially with the number of units in the homeostat.

**Both:** More questions and ideas about these and other AI problems.

**Ray's notes:** Ashby is interested in bringing some rather important simple ideas to the psychologists... Ashby "is fascinated by the fact that a table of random numbers can give one all one wants — i.e. that selection is the only problem."

McCarthy's view, that the only real problem is the search problem — how to speed it up. Ashby feels the same but has fewer ideas. Ashby wants to start slowly, like organic evolution, and work up to more complex phenomena



Pictures of a homeostat as sketched by Ray during Ashby's talk and later drawn carefully by Trenchard.

### Samples of other discussions:

**Trenchard, July 24:** Described different levels of complexity in the machine, the program and the language used, in order to develop a deductive proof for axioms. He stressed different kinds of symbols, what they would connote and what associations they might have, and discussed a theory of machine semiology, or how signs could be applied to machine design, and how machines could understand signs.[21, pp. 5]

**Samuel, Aug 6:** talked about checkers: "Board positions — how to introduce intermediate goals. How to score different types of moves."<sup>3</sup> ... "Samuel is entirely interested in the learning problem, not so much in checkers."

**Bernstein, August 8:** on chess, explicating a player's intuitive ideas, such as: "for learning: try to re-weight so that the winning player's moves would

<sup>3</sup>"About 1/2 the time is spent on position evaluation, 1/2 on updating. It turned out that the good evaluation method wasn't so important, as it played well when the evaluator was 'seriously wrong'!"

have been made. Looks at losing opponent's play to see in which way his moves appear to be weak."

**Selfridge, no date given:** Interested in the learning process. What sentences are close to a question sentence. What is reasonable logic for a game; examples of useful words in the game.

**Julian Bigelow, Aug 15:** "Very wary of speaking vaguely with the hope of being able to translate this vague talk into machine language (McCarthy also)."<sup>4</sup> (The foregoing are sampled from Ray's more detailed notes.[37])

**More discussions — McCarthy, Minsky, Simon:** led discussions, some referring to recent and current papers such as McCarthy's "On the Meaning of 'The Reader will Easily Verify'," Minsky's "A Framework for Artificial Intelligence," and Newell, Shaw and Simon's "Logic Theorist."

### Did They Find a Group Project?

Ray's notes during McCarthy's talk, August 7, indicate that McCarthy sounded out the group on their interest in chess, possibly using the IBM 704. Ray noted their reactions:

**McCarthy:** — "still strong for chess"

**Bigelow:** — "tolerant of chess, and any other approach, is not so optimistic about how good T.M.'s can be in a few years"

**Shannon:** — "interested in chess, but isn't too enthused over its relevance to the art. int. problem"

**Selfridge:** — "doesn't think chess is so great, but will go along with McC"

**Samuel:** — thought checkers would be as good though he "doesn't seem too happy about how to get learning out of checkers, either"

**Nat Rochester:** — preferred to solve problems from Feller's textbook on probability — liked Marvin's ideas but those weren't yet concrete enough to start on — "isn't too hot on chess, — but will go along with McC if there is nothing better"

**Minsky:** — "*very* much against chess, also would like to get more active interest in his own approach, which is Geometry, with figure drawing and a kind of model"[33]

Did this inaugurate a group project? No.

### What Else Did They Do?

They talked informally: Trenchard reminisced, "The person that impressed me most was Abraham Robinson. I can remember walking the golf course with Abraham Robinson and a couple of others and not understanding a word that he was saying. But at that time he was becoming probably the best mathematician in the world." [23]

---

<sup>4</sup>"On the other hand, Minsky, Rochester, and perhaps More are less wary of this, and Minsky in particular is not afraid to describe things at a fairly intuitive level, then become more and more concrete."



They had afternoon tea; a note written in 2002 by Trenchard describes an incident with the mathematician John Nash: “He and I were following Herb Simon, Allen Newell, and others up the stairs to an apartment for late afternoon tea, when Nash commanded sharply: “Pick that up!” I glanced at the odds and ends on the landing and said: “Pick it up yourself!” He replied that he was testing me.”[22]

During these weeks Ray wrote 175 pages on induction, describing matrices of symbols which could be program code, that predicted future matrices. He wrote it up in a paper, “An Inductive Inference Machine,” Ray’s first report on induction for intelligent machines. He gave two talks on it.[34]

### Did They Have The Tools They Needed?

AI was just beginning to get vital tools. The earliest programming languages were written by some participants. Newell, Shaw and Simon had already developed a series of languages called IPL (information-processing language) for their work. In 1958 McCarthy developed LISP which, being easier to use, replaced IPL.[26] At the AI@50 conference, Oliver Selfridge told how essential the development of improved languages and machines was, and described early pioneers such as J. C. R. Lickleiter developing time-sharing, Nat Rochester designing IBM computers, and Frank Rosenblatt working with perceptrons.[19]

In fact we have Trenchard’s notes only because Nat Rochester was so busy working on the IBM 704 — that room-sized computer using vacuum tubes — that he sent his student Trenchard to attend three of the weeks in his place.[23]

And Trenchard remembers the day one of his colleagues come in at MIT to demonstrate a brand-new tool: the computer mouse![23]



Trenchard and his wife Kit in 2011 and Trenchard at work on his Lattice Array Theory in 2011.

The summer project had other practical difficulties: Later John McCarthy said “You see my intention when I first thought ... about organizing it — The whole group would meet for the whole summer. Two things prevented that; one was that Rockefeller didn’t give us anywhere near enough money for that; and the people had other commitments and so they came for short times and different times.”[14] Many participants only showed up for a day or even less.

## Was It A Success?

Glowing words have called the Dartmouth Summer “The Birth of Artificial Intelligence,” introducing the major AI figures to each other.[28, p. 17]

Both Trenchard and Ray gave a summary in their notes. Trenchard was positive, detailed the specific interests of each member and felt he had been influenced by many members.[21, pp. 2–19]

Ray’s comment was austere:

Th. research project wasn’t very suggestive. The main things of value:

- 1) wrote and got report reproduced (very important)
- 2) Met some interesting people in this field
- 3) Got idea of how poor most thought in this field is
- 4) Some ideas:
  - a) Search problem may be imp.
  - b) These guys may eventually invent a T.M. simply by working more and more interesting special problems.[39, p. 2]

So the project never created a clear direction toward the goal of AI. But several results came out of the summer that had lasting effects.

## Part Two: Ideas that Mattered: Three Workshop results and what came after

Few papers came from the summer. McCarthy intended to have several speakers and reports, including Ray’s paper, at the September 1956 IRE Symposium on Information Theory. Unfortunately, he missed a deadline, so only “Logic Theorist” by Newell and Simon and a paper by Rochester et al. on computer modeling of Hebb’s cell assemblies were included in the proceedings. McCarthy gave a summary of the summer, and Ray circulated his paper on the inductive inference machine. Neither McCarthy’s summary or Ray’s paper made it into the proceedings.[29]

McCarthy sent Ray’s paper to Morison in November 1956, who replied “We are naturally pleased to find that the seminar progressed far enough so that some of the thoughts presented could be worked out so fully.”[25]

The participants had so many ideas and such varied projects, that perhaps the summer’s main function was to inspire, helping them understand and develop their own work better.

Trenchard, for example, mentions many who influenced him, including Shannon, McCarthy, Rochester, Minsky, Solomonoff, Selfridge, Bigelow, Robinson, McCulloch, Newell and Simon.[31] Trenchard’s own work led him later to create array theory and in fact, his own version of a “Turing Test” for AI.

But there were three particular aspects relevant to AI that this report will discuss. Because of my interest and knowledge of Ray’s work, the focus is on how these aspects relate to Ray. I can’t do justice to the impressive work of the

other participants, though I wish I knew enough to do so. But Ray can serve as a focus to describe these three aspects.

## 1. Semantic, Numerical and Symbolic Information

### Non-Symbolic Methods

Researchers consider a main result of the Dartmouth Workshop to be the debut of the “new paradigm” of symbolic information processing, represented by Newell and Simon’s ideas,[16] replacing the earlier methods of brain modelling associated with cybernetics and neural nets.[8, p. 5]

Norbert Wiener had defined cybernetics to be the science of control and communication in the animal and machine,[8, p. 5] deriving the word “Cybernetics” from the Greek for “Steersman.”[47, p.9] Brain modelling, the close analogy between the activity of neurons and the digital activity of computers, was explored by many researchers, such as Ashby and Milner and Minsky.

Minsky had been using semantic methods for problem solving, trying to apply this to learning by neural nets. After the Dartmouth Summer, Marvin’s ideas changed from an interest in brain modelling to symbolic AI.[8, pp. 5, 12, 13]

### Symbolic Methods

During the summer Marvin had many discussions with Ray about his inductive inference paper on how machines could be made to improve themselves using symbolic methods with associated utility values of how successful they were at prediction.[40] Even at Ray’s first talk, he discussed the advantages of using symbols for induction.[32, p. 3]<sup>5</sup>

### Newell, Shaw, and Simon

Allen Newell, J.C. Shaw, and Herb Simon had been working on a report on symbolic logic for many months. Their paper, “Logic Theorist,” finished by the summer’s end, became a forerunner of early AI programs. It focused on processing symbol structures, using heuristics as being fundamental to intelligent problem solving.[26, p. 54] They completed it in time for the 1956 IRE Symposium.

Simon, and probably Newell, saw Ray’s report, but were not influenced by it. Simon later did recall that Ray’s paper (version of Aug. 14, 1956) was discussed there: “a very imaginative proposal for a system that would learn (all sorts of things) from worked-out examples.”[30]

---

<sup>5</sup>This lecture, on July 10, was attended by Minsky, McCarthy, McCulloch, Shannon, and Bill Shutz. Some reactions were: Shannon: do a hand simulation of this; McCarthy: worried about length of some of the potential search processes; Minsky: happy to add many ad-hoc mechanisms if useful (which Ray much disagreed with).[34, pp. 8–10] On the last day of the workshop, Ray gave another talk on the paper, one or more versions of which he had circulated at the summer session.

## Transitions toward Symbols


Others were also interested in symbolic language. An example is Trenchard's work. In his June 12, 1956 paper "Computer Decisions in Deductive Logic"[20] he discusses the use of symbols being more general than arithmetical methods, and using deductive methods rather than lookup decision trees. Trenchard later expanded his ideas into the creation of Array Theory, which led to the understanding and better use of arrays. (Arrays can even be empty. And an empty array, empty of 2, is different from an empty array, empty of 3! A mystery!)

## Minsky and Ray

But Marvin was most influenced by Ray's work in symbolic methods.

In Ray's memorial pages, he writes: "In the early 1950s I developed many theories about ways to make neural networks learn — but none of these seemed likely to lead to high-level kinds of idea. But once Ray showed me his early ideas about symbolic "inductive inference scheme" these completely replaced my older approach. So Ray's ideas provided the start of many later steps toward Artificial Intelligence."[5, p. 1]

In an interview with the author Pamela McCorduck, Minsky said, "the explanation of my perhaps surprisingly casual acceptance of the Newell-Shaw-Simon work — ["Logic Theorist"] was that I had sketched out the heuristic search procedure for [a] geometry machine and then been able to hand-simulate it in an hour or so..."[16, p. 126] In another talk with me, Minsky noted that the influence toward the change in his ideas toward symbolic logic stemmed from Ray's work. Simon and Newell did not know this, which probably was the cause of some friction.



Marvin playing the recorder at Ray's apartment in Harvard Square.

## Later Developments

Newell and Simon were by far the most influential Dartmouth Summer attendees contributing to the AI community's shift to symbolic information processing, having helped originate it, continuing with "Logic Theorist," and later developing new symbol-manipulation strategies in "General Problem Solver;" eventually they received the prestigious Turing Award for their work.[26, pp. 81, 88] Primarily through Marvin, Ray also was influential. But since no version of his Dartmouth paper was published until 1957, since he was not affiliated with any university, and since he was focused on induction for prediction, rather than on deduction strategies, his contribution wasn't noticed — not even by Ray, himself!

## 2. "Deep" versus "shallow" methods

In his August 7 talk, McCarthy spoke about going deep into a specific problem.<sup>6</sup> Elsewhere Ray's notes point out: "if he [McCarthy] used a particular system [Smulyan's logic rules] then all the information the machine could gain would be information about these Smulyan formal systems. It is improbable that such rules about program improvement would be useful in, say, solving Trig. identities." [35, pp. 8–9]

Ray wanted to develop a very general system: "Idea of working on many humanly solvable problems lightly, to find out what they have in common, rather than initial intense work on any one problem. Emphasize how human works problem rather than easy trick ways to work problems that may be peculiar to that set of problems." There has been "a great overemphasis of difficulties peculiar to the particular set of problems worked on rather than difficulties peculiar to all problems. Use of computers is, I think, usually going in more deeply than is good at the present time." [35, p. 8]

In discussing a general common-sense program rather than a specialized program delving deeply into one field, Ray mentioned considerations such as: in a Math TM (thinking machine) "absolutely rite or wrong answers do not correspond to what is met in th. ordinary language." [32, p. 14]

## Limited-Domain Projects

However most participants — Ashby, More, Samuel, McCarthy, Bernstein, Minsky, Bigelow — used more limited and well-defined domains and problems.<sup>7</sup> Some examples:

---

<sup>6</sup>e.g., "[In a board game] there is an optimal *constant* time for deciding how good a position is. Therefore if we play *very* deep, its time should be negligible compared to the amt. of time spent updating positions." McCarthy's updater, the "plausible move program," has a special operator with complex procedures.

<sup>7</sup>When Marvin Minsky wrote about building a logic-proving machine (July 6, 1956): "Actually, the machine is intended to operate over much broader classes of problems, including the making of empirical perceptions and judgements." Ray noted in the margin that it sounded "suitable mainly for well-defined problems" [17].



**Marvin Minsky’s geometry machine:** Quoted from Ray’s notes on Marvin’s discussion, August 13 —

Make machine to prove theorems in ordinary plane geometry. Statements of theorem are in modified English. The machine makes up proofs by drawing diagrams to suggest various sub-goal theorems. . . . Basically, I think his primary idea is to study heuristic devices — Eventually, to make generalizations about them, so that one can get *methods* of inventing new ones.[] (boxaaugust 13)

Rochester probably attended that session; certainly he talked a lot with Minsky about Marvin’s geometry machine ideas and became very enthusiastic, describing these ideas afterward to Herb Gelernter, who went on to research a geometry-theorem-proving machine, presenting a major paper on this in 1959.[26, p. 85]

**Arthur Samuel’s checkers program:** He had a continuing research project on a checkers-playing program using several different representations for checkers knowledge, and different training methods. Eventually he developed a program that played better than he did! In 1965, however, the world champion, W.F. Hellman, beat it in all four correspondence games and only drew one other game because, as Hellman explained, it was a “hurriedly played cross-board game.”[28, p. 18] [3, pp. 457–463]

### The Demo to Sponsor Problem

Later, Ray notes that Wendy Conquest, (a creator of the movie “Mind in the Machine: The Discovery of Artificial Intelligence”) “felt that I was a strong proponent of doing very easy problems as opposed to most of the rest that wanted to do hard problems to show that their techniques were actually useful (The “demo to sponsor” problem).”[46]

The Dartmouth Summer itself had a “demo to sponsor” problem. The Rockefeller Foundation gave about half the requested funding. Robert Morison, who was director of Biological and Medical Sciences at Rockefeller, wrote: “I hope you won’t feel that we are being overcautious, but the general feeling here is that this new field of mathematical models for thought, though very challenging for the long run, is still difficult to grasp very clearly. This suggests a modest gamble for exploring a new approach, but there is a great deal of hesitancy about risking any very substantial amount at this stage.”[25]

Early impressive programs assisted in a particular field: DENDRAL to assist in chemistry; MACSYMA in mathematics; the Sussman-Stallman program for understanding electronic circuits;[16, p. 318] These became known as Expert Systems. They used procedural knowledge, solving problems in specific situations using procedures that were similar to the way people explain how they solve problems.[1, p. 198]

In 1985 Ray wrote, “. . . commercial success of AI — mainly Expert Systems — has lured many bright graduate students away from general theory, to work on industrial applications.”[44, pp. 149–153]

And as the cold war got frigid, in 1969 Congress passed the “Mansfield Amendment” to the Defense Procurement Act, demanding that basic research funded by ARPA (soon renamed DARPA to stress defense) had to have “a direct and apparent relationship to a specific military function or operation.” Wonderful civilian research, such as that done at BBN, got choked off! In the following years, a continuing series of Mansfield Amendments further narrowing appropriations from the military, channeled research still more.[26, p. 20]

### 3. Deductive logic versus inductive probabilistic methods

Another importance of the summer: The trend toward machines solving specialized problems rather than acquiring general learning was accompanied by a nascent split between deductive, logic-based methods and inductive, probabilistic methods.

#### Two Aspects of Thinking

Here are two major categories of reasoning: **Deductive Reasoning** draws logically valid conclusions from some assumed or given premise. We use this for things like math proofs or in dealing with formal systems. But human thinking also includes predictions that involve estimations and generalizations. For this we use **Inductive Reasoning** — drawing the “best” conclusions from a set of observations. We create a theory to explain past data, and observe how well the theory works on new data.[27, p. 1077]

Some workshop participants considered probabilistic methods, but most focused on deductive logic projects. Ray was the only member of the group who focussed completely on induction using probabilistic measures, and his understanding of probability was still in an early stage.

#### A Little Interest in Probabilistic Reasoning

Some comments suggested the usefulness of probabilistic induction:

**Trenchard’s notes:** “Warren McCulloch on July 25, discussed Carnap and Turing machines [A theoretic specialized computer now used in probabilistic prediction]. He believed the human brain is a Turing machine.” — “A Turing machine with a random element in the control unit, is the same as a Turing machine with a random tape . . .”[21, p. 4]

**Ray’s notes:** Selfridge at one point stated that “deductive logic is rather useless” for certain machines . . . and “also S. thinks some statistics are necessary for learning — contrary to McC.”

Nat Rochester was interested in using AI to solve probability problems in a book by Feller, which Ray thought was too hard to start on.[33, p. 3]<sup>8</sup>

---

<sup>8</sup>This is the classical text “Probability Theory and Its Applications,” by William Feller, Eugene Higgins Professor of Math at Princeton University. In one section relevant to Ray’s work, it derives formulas of how to update the likelihood of some events as the sample spaces including them are updated (Bayes’ rule); it states that “If the events . . . are called causes, then [the formula] becomes Bayes’ rule for the probability of causes.”[4, p. 85]

Participants also felt that induction would eventually be very important, and that deductive procedures could be the gateway toward this goal. For example, Trenchard noted . . . “concepts of this program [in his work proposed] may then be extended to . . . mathematical induction” . . .

### A Lot of Interest in Deductive Reasoning

Most work by the participants used deduction and heuristics, such as Newell, Shaw and Simon’s “Logic Theorist,” which rearranged symbol structures toward a certain goal, using heuristics to guide the rearranging. Ray noted that McCarthy “felt using probabilistic methods would blossom into too much space . . . the problem of search,” but Ray described some of the problems with thinking machines using non-probabilistic methods — it could “be seriously disturbed forever by one ill-conceived example.” — while a probabilistic machine would be “not *too* much disturbed by a single counter example.”[36, p. 8]

Later McCarthy said:

I was focusing on expressing common sense knowledge in formulas of mathematical logic. But not until '58 really. . . I was sort of working up to it in '56 to get facts by observation and then use them in connection with these general rules to determine that a certain strategy will succeed in achieving a goal.[14]

He planned ordered steps and a strategy, keeping away from statistics since success/failure decisions would get too large too fast and the criteria for success too removed from the elementary act.[12]<sup>9</sup>

### McCarthy and Ray

Amazingly, in spite of being quite negative about probabilistic methods, McCarthy significantly influenced Ray during the summer. The clearest description was written later by Ray, in 1997:

“One day McCarthy gave a talk on “Well Defined Mathematical Problems.”<sup>10</sup> . . . I asked him about the induction problem: ‘Suppose you are given a long sequence of symbols describing events in the real world. How can you extrapolate that sequence?’ Next day he said, ‘Suppose we were wandering about in an old house, and we suddenly opened a door to a room and in that room was a computer that was printing out your sequence. Eventually it came to the end

---

<sup>9</sup>McCarthy presented these ideas in his later paper on well-defined problems, “Programs with Common Sense,” at the symposium on “Mechanization of the Thought Process” in Teddington England, 1958.[13] Later he developed a formula called “circumscription” which is closer to probabilistic induction: “minimum principles in logic is that a predicate is true for a minimal set of objects, so it’s the set of things which is true which can’t be reduced further. That’s what the circumscription formula says.”[14]

<sup>10</sup>His thesis was that all mathematical problems could be formulated as problems of inverting Turing machines. Specifically, given a machine  $M$  whose inputs were strings of symbols. Given a desired output string,  $s$ , we are required to find a string  $p$  such that  $M(p) = s$ . McCarthy gave many examples to show how problems could be put into this form.[45]

of the sequence and was about to print the next symbol. Wouldn't you bet that it would be correct?' ... I remembered the idea because it seemed intuitively reasonable.”[45]<sup>11</sup>



McCarthy around 1956.

In part, McCarthy's ideas at Dartmouth helped Ray reach his major work, which became a foundation for algorithmic information theory: algorithmic probability and his general theory of inductive inference. This is a mathematical system for predicting future theories based on weights related to the lengths of present theories coded as strings, with the shortest having the most weight.[42][43]

Since that time, probability has become well established in AI, and probabilistic methods are used in AI.

But the early focus on deductive methods, especially the popularity of rule-based expert systems, led to probability being scorned for a while. What is now called AI often focuses on logic-based, and “deep” systems. Inductive and probabilistic methods are now often categorized as Machine Learning, holding separate conferences.

Logic-based systems still don't generalize very well and lack superior methods of learning.

Prediction by induction has not achieved the AI goals either, but several projects are moving in that direction, such as the work of Marcus Hutter, on a system called AIXI which has a reinforcement learning agent that uses a form of algorithmic probability to weight programs and the rewards they generate depending on the agent's next action.[9]

### **Marvin says nice things about Ray, still prefers GOFAI<sup>12</sup>**

In a round table discussion in 2011, with Marvin, Gloria Minsky, Peter Gacs, Ed Fredkin, Cynthia Solomon, Alex Solomonoff, Charlie Bennet (on Skype),

---

<sup>11</sup>Ray wrote a long note about this in the July 23 1956 meeting: saying “This idea may, indeed, work.”[38]


<sup>12</sup>Good Old Fashioned AI

and myself we talked about the Dartmouth paper. Marvin said

“Anyway this is the paper [Ray’s paper at Dartmouth] and it’s full of these little N-grams and I was very impressed by this and this led me to switch from working on neural nets to working on symbolic ideas about AI. Neural nets just learned by computing correlations between things and sort of in a way it was a kind of simple statistical type of learning whereas Ray’s idea was you would consider the possible meanings of different rearrangement of symbols and from my point of view the progress of artificial intelligence in the 60s and 70s was rapid and very impressive because everybody that I knew was working on symbolic reasoning and representation and that sort of thing, and then pretty much starting in the 1980s people started to go back to statistical learning so that the artificial intelligence being done now resembles the kind that was being done from 1900 to 1920 when people like Watson and Pavlov were doing statistical learning. So in some sense I see Ray as having started the modern direction, then when he made this more beautiful mathematical theory it seems to me that the general theory of induction that he produced solves lots of philosophical problems and it sort of ended 3000 years of speculations in philosophy of what is meaning and what is information and so forth. But I’m not sure that it’s a good line to pursue.” Marvin concludes, “As a tour de force we should republish the old paper ‘The Inductive Inference Machine’ ”[(marvingroup1)

### Probability Gets Applause

But others feel positive about probabilistic learning. At the AI@50 conference, the debate continued about whether AI should be logic-based or probability-based.[19, p. 88] “In defense of probability David Mumford argued that in the last 50 years, brittle logic has been displaced with probabilistic methods. Eugene Charniak supported this position by explaining how natural language processing is now statistical natural language processing. He stated frankly, ‘Statistics has taken over natural language processing because it works.’”(Quote by James Moor)[19, p. 89]



Ray at the AI@50 conference.



## It Was Fun!

These are some examples of what the Dartmouth Summer was like, how it began, what happened during those summer weeks and after. The summer conference was influential, in ways less obvious than the few well-known programs that came out of it. Several directions were initiated or encouraged by the Workshop, debated to this day. By understanding what each other was thinking, and why, the participants got valuable new ideas, even when disagreeing. The most interesting and direct changes with respect to Ray, as described here, were the change in Marvin's ideas, due to Ray; and a development of Ray's ideas, due to McCarthy.

There were many other inspired ideas, such as Selfridge's thoughts that later led to his "Pandemonium", Rochester's work on cell assembly theory, More's work leading to his "NIAL" programming, Marvin's paper (July 5, 1956), "A Framework for Artificial Intelligence" which foreshadowed his seminal work "Steps toward Artificial Intelligence".

Regardless of what papers and philosophical views emerged, the Dartmouth Summer of 1956 was a wonderful event. A group of amazing people got together and freely exchanged ideas and enthusiasms. As Arthur Samuel said, "It was very interesting, very stimulating, very exciting." [16] It was not a directed group research project. It was more like inviting a bunch of brilliant people to an eight-week conference party, where everyone was brimming over with their own ideas. Ray always remembered that time as part of the joy and promise of those beginning days of AI.

AI has a colorful and contentious history. Many scientists and mathematicians have contributed to it and been influenced by each other in meaningful but little-known ways. Ray was there at one of AI's bright early moments: the Dartmouth Summer Workshop of 1956. Hopefully this paper, by looking at what happened, through early notes, has added some new insights into that summer, and how it affected Ray's work as well that of others. Papers of other scientists yet to be discovered will add to the wealth.



The five attendees at the AI@50 conference in 2006 who were part of the Dartmouth Summer of 1956: Trenchard More, John McCarthy, Marvin Minsky, Oliver Selfridge, Ray Solomonoff.

**Further Reading:** For more scholarly discussion of other scientists at Dartmouth, and what followed, Nils Nilsson’s book, “The Quest for Artificial Intelligence” is an excellent source, as is Pamela McCorduck’s “Machines Who Think.”

**Acknowledgements:** planning to use part of thanks. . . other acknowledgements. . .

All the Dartmouth documents and recordings from my files that are referenced in this paper are or will be available on Ray’s Memorial website at <http://raysolomonoff.com/dartmouth/dart.html> The link for Ray’s paper from that summer is on his publications page at: [raysolomonoff.com/publications/indinf.pdf](http://raysolomonoff.com/publications/indinf.pdf).)

## References

- [1] A. Barr and E.A. Feigenbaum, editors. *the Handbook of Artificial Intelligence*, volume 1. William Kaufmann, Inc, 1981.
- [2] R. Carnap. *Logical Foundations of Probability*. 1950.
- [3] P. Cohen and E.A. Feigenbaum, editors. *the Handbook of Artificial Intelligence*, volume 3. William Kaufmann, Inc, 1982.
- [4] W. Feller. *Probability Theory and Its Applications*, volume 1. Wiley and Sons, 1950.
- [5] Memorial Guests. *Ray Solomonoff Memorial Notes*. March 2010.
- [6] McCarthy J. interestalist. to be published on [raysolomonoff.com](http://raysolomonoff.com), Sept 1956.

- [7] Kirby. kirletter.pdf. to be published at <http://raysolomonoff.com/dartmouth/dart/misc>, Nov 2011.
- [8] Ronald R. Kline. Cybernetics, automata studies and the dartmouth conference on artificial intelligence. *IEEE Annals of the History of Computing*, October-December 2011. Published by IEEE Computer Society.
- [9] Hutter M. *Universal Artificial Intelligence*. Springer-Verlag, 2004.
- [10] J. McCarthy. Mccarthy\_to\_morison\_5\_25\_56.pdf. Rockefeller Foundation Archives, Dartmouth file, 1956.
- [11] J. McCarthy. mccarthy\_2\_23\_56\_letter\_mentions\_invite\_for\_newell\_and\_simon.pdf. Rockefeller Foundation Archives, Dartmouth file, February 1956.
- [12] John McCarthy. Dartmouth mathematics project. December 1956. dartmouth\_mathematics\_project.pdf; a privately circulated report.
- [13] John McCarthy. Programs with common sense. In D.V. Blake and A.M. Uttley, editors, *Proceedings of the Symposium on Mechanization of the Thought Process*, London:Her Majesty's Stationary Office, 1959.
- [14] John McCarthy. Interview with mccarthy. To be published in raysolomonoff.com, 2011.
- [15] Minsky M. Rochester N. Shannon C.E. McCarthy, J. A proposal for the dartmouth summer research project on artificial intelligence (dart564props.pdf). Box A, August 1955.
- [16] P. McCorduck. *Machines Who Think*. A.K.Peters, Ltd., second edition, 2004.
- [17] M. Minsky. A framework for an artificial intelligence. July 1956. A privately circulated report.
- [18] M Minsky. minsky\_3\_9\_56\_letter\_to\_john\_nat\_and\_claude.pdf. Rockefeller Foundation Archives, Dartmouth file, July-August 1956.
- [19] J Moor. The dartmouth college artificial intelligence conference: The next fifty years. *AI Magazine*, 27(4):87–91, 2006.
- [20] Trenchard More. Computer decisions in deductive logic (moreprop.pdf). Box A, June 1956. report from Dartmouth Summer AI Project.
- [21] Trenchard More. dart56more5th6thweeks.pdf. <http://raysolomonoff.com/dartmouth/dart/boxa>, summer 1956. typed notes.
- [22] Trenchard More. trenattend.pdf. <http://raysolomonoff.com/dartmouth/dart/boxa>, summer 1956. typed notes.

- [23] Trenchard More. home interview. To be published in raysolomonoff.com, September 2011.
- [24] Morison. [Rockefeller\\_to\\_shannon\\_and\\_mccarthy\\_6\\_17\\_55.pdf](#). Rockefeller Foundation Archives, Dartmouth file, June 1955.
- [25] R Morison. [Morison\\_to\\_mccarthy\\_11\\_21\\_56.pdf](#). Rockefeller Foundation Archives, Dartmouth file, November 1955.
- [26] N. Nilsson. *The Quest for Artificial Intelligence*. Cambridge University Press, 2010.
- [27] S. Rathmanner and M. Hutter. A philosophical treatise of universal induction. *Entropy*, 13:1076–1136, 2011.
- [28] S. Russell and P. Norvig. *Artificial Intelligence, A Modern Approach*. Pearson Education, second edition, 2003.
- [29] H Simon. [3\\_of\\_11.pdf](#). Rockefeller Foundation Archives, Dartmouth file, Nov 1999.
- [30] H Simon. [5\\_of\\_11.pdf](#). Rockefeller Foundation Archives, Dartmouth file, Nov 1999.
- [31] H Simon. [9\\_of\\_11.pdf](#). Rockefeller Foundation Archives, Dartmouth file, Nov 1999.
- [32] R.J. Solomonoff. [dart1n71856\\_0115.pdf](#). <http://raysolomonoff.com/dartmouth/dart/notebook>, July-August 1956.
- [33] R.J. Solomonoff. [dart56aug7chess.pdf](#). <http://raysolomonoff.com/dartmouth/dart/boxa>, August 1956.
- [34] R.J. Solomonoff. [dart56ray622716talk710](#). <http://raysolomonoff.com/dartmouth/dart/boxbdart>, July-August 1956.
- [35] R.J. Solomonoff. [dart56rayn\\_notesonothers.pdf](#). <http://raysolomonoff.com/dartmouth/dart/boxa>, August 1956.
- [36] R.J. Solomonoff. [dart7n81256\\_89102.pdf](#). <http://raysolomonoff.com/dartmouth/dart/notebook>, August 1956.
- [37] R.J. Solomonoff. [dartmynotes723815.pdf](#). <http://raysolomonoff.com/dartmouth/dart/boxa>, August 1956.
- [38] R.J. Solomonoff. [dartmynotes723815.pdf](#). <http://raysolomonoff.com/dartmouth/dart/boxa>, July-August 1956.
- [39] R.J. Solomonoff. [dartwho\\_overallsummary.pdf](#). <http://raysolomonoff.com/dartmouth/dart/notebook>, August 1956.

- [40] R.J. Solomonoff. An inductive inference machine. Dartmouth Summer Research Project on Artificial Intelligence, August 1956. A privately circulated report.
- [41] R.J. Solomonoff. rayaiapproach.pdf. <http://raysolomonoff.com/dartmouth/dart/boxa>, January 1956.
- [42] R.J. Solomonoff. A formal theory of inductive inference: Part I. *Information and Control*, 7(1):1–22, March 1964.
- [43] R.J. Solomonoff. A formal theory of inductive inference: Part II. *Information and Control*, 7(2):224–254, June 1964.
- [44] R.J. Solomonoff. The time scale of artificial intelligence; reflections on social effects. *Human Systems Management*, 5:149–153, 1985.
- [45] R.J. Solomonoff. The discovery of algorithmic probability. *Journal of Computer and System Sciences*, 55(1):73–88, August 1997.
- [46] R.J. Solomonoff. raywendypoints.pdf. <http://raysolomonoff.com/dartmouth/dart/boxa>, August 2005.
- [47] N. Wiener. *The Human Use of Human Beings*. Houghton Mifflin Co., 1950.
- [48] C. Wills. *America in the 1950s*. Stonesong Press, 2006.