Early Artificial Intelligence

Early artificial intelligence was a mix of ambition and self-discipline. Some scientists were quick to over promise what could be done with early machines. At the same time, you could see the potential in these machines to solve complex problems. In 1956, you had one of the first attempts to create a machine with general intelligence. Allen Newell and Herbert Simon created a computer program they called the General Problem Solver. This program was designed to solve any problem that could be presented as mathematical formulas. One of the key parts of the General Problem Solver was what Newell and Simon called the physical symbol system hypothesis.

In their paper, they argued that symbols were the key to general intelligence. If you could get a program to connect enough of these symbols, then you would have an intelligent machine. Symbols are a very big part of how you interact with the world. When you see a stop sign, you know to look for traffic. When you see the letter a, you know that the word will make a certain sound. When you see a sandwich, you might think of eating. Newell and Simon argued that if a machine were trained to understand these symbols, they could behave more like humans.

They thought a fundamental part of human reasoning was just connecting these different symbols. In one sense, our language, ideas, and concepts were just broad groupings of interconnected symbols, but not everyone bought into this idea. In 1980, philosopher John Sorrell debated that you could never call these symbolic connections real intelligence. To explain, he created something called the Chinese room argument. In this experiment, you should imagine yourself in a windowless room with one narrow slot on the door, almost like a mail slot.

You can use the slot to communicate with the outside world. In the room you have a book on a desk, and bunch of Chinese symbols on the floor. This book is filled with long lists of matching patterns. It says if you see this sequence of Chinese symbols, then respond with that sequence of Chinese symbols. To begin the experiment, John Sorrell imagines that someone who speaks fluent Chinese writes a note and shoves it through the slot. You pick up the note, but you don't speak a word of Chinese, so you really have no idea what it says.

Instead, you simply go through the tedious process of looking through your book and matching the sequence of Chinese letters. Once you find the matching sequence of letters, you look at how the book tells you how to respond. Then you tape together the reply from the Chinese symbols and push your note through the same slot on the door. A native Chinese speaker on the other side might believe that they're having a conversation. In fact, they may even start

assuming that the person in the room is intelligent, but Sorrell debates that this is far from intelligence since the person in the room can't speak Chinese and doesn't have any idea what you're talking about.

It was just simply matching patterns. You can try a similar experiment with your smart phone. If you ask Siri or Cortana how they're feeling, they'll give you a response. They'll usually say they feel fine, but that doesn't mean that they really feel fine. In reality, they also don't really know what you're asking, they're just matching your question to a pre-programmed response just like the person in the Chinese room. So Sorrell debates that simply matching symbols is not a true path to intelligence, that a computer is acting just like the person in the room.

They don't understand the meaning or the content. They're just matching symbols from a long list of instructions. You can also see how the book of Chinese responses might become larger and larger as you try to create more matching statements. This is called combinatorial explosion. There are so many different combinations that matching becomes overwhelming. Think about all the things that people might ask in Chinese, or imagine all the different responses that Siri must be prepared to answer.

Even with these challenges, physical symbol systems were still the cornerstone of AI for 25 years, yet in the end creating all these matching connections took up too much time. It was also challenging for these machines to match all the different patterns without running into combinatorial explosion.