Temperature and Entropy in AI

By: Jimmy Shah

What is AI?

It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

-John McCarthy

Artificial intelligence, or AI, is the field that studies the synthesis and analysis of computational agents that act intelligently.

- Poole & Mackworth

In practice, AI is about making computers and other machines perform tasks that (for humans) seem to require intelligence

-Bryce Wiedenbeck, CS Department

What AI is NOT (yet)

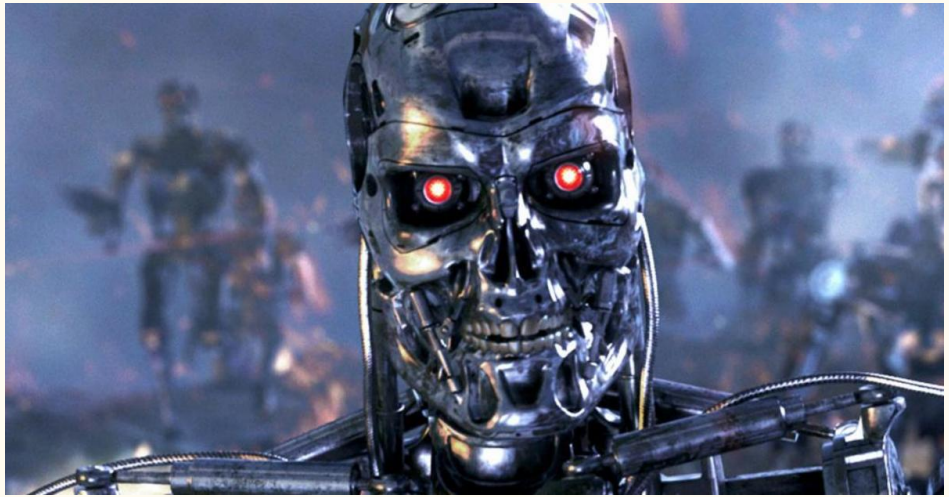

Searching...

One application of AI -> Searching!

Using “intelligence/heuristics” to make better moves without exhaustive search

Traveling Salesman Problem:

> *"Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city and returns to the origin city?”*

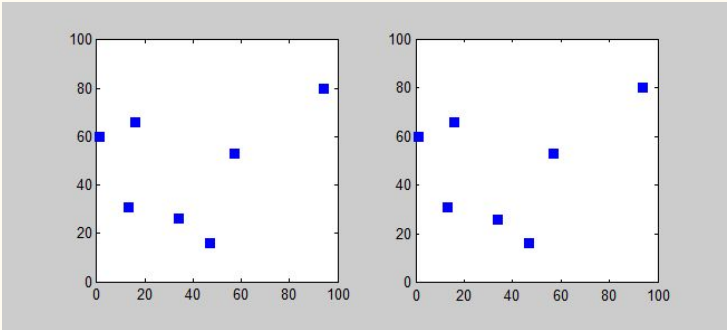

Hill Climbing

Key idea:

- Start with an arbitrary candidate.
- Iteratively move to a neighbor with higher value.
  - Climb up the value hill.

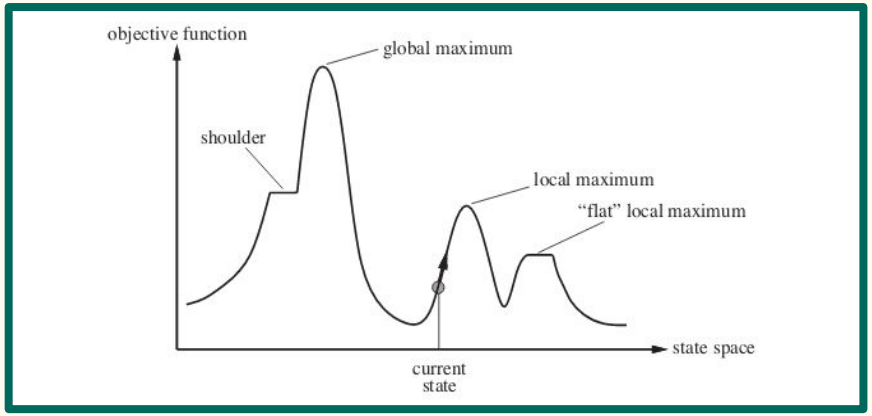

Branching Factor:

- Searches can be slow, if states have lots of neighbors
- Alternatively, pick a random neighbor. If better, keep!
  - Else, try again

Problems:

- We can get stuck in local optima
- Possible Solutions:
  - Random steps
  - Random restarts

Use Temperature To Help Search!!!

Simulated Annealing

Big Idea:

- combines hill climbing with a random walk
- Select moves randomly, but accept bad moves with low probability!

Setup:

- Have an initial temperature
- Set a decay rate

$$Prob = e^{\frac{cost_{state} - cost_{neighbor}}{t * rate_{decay}^{rounds}}}$$

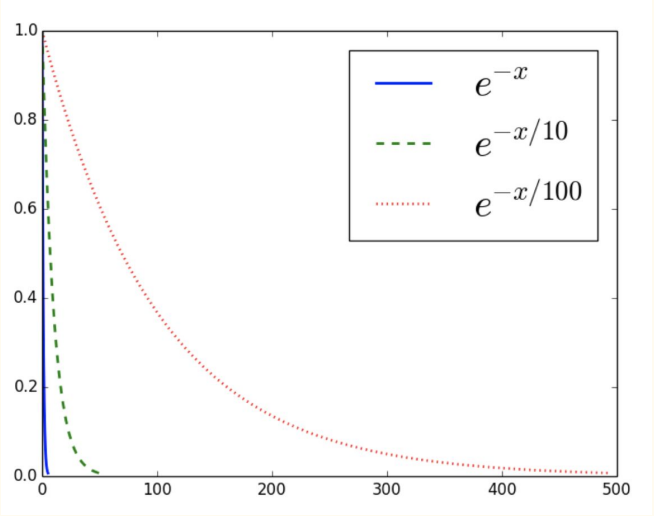

Stochastic Beam Search

Rather than random restarts, want to focus on promising paths

Could do a parallel search, but we lose good candidates quickly

Idea:

Use Gibbs Sampling to select a new population

- Probability similar to SA
- Have each candidate have a weight
- Create a distribution from weights
    - Then set population from distribution

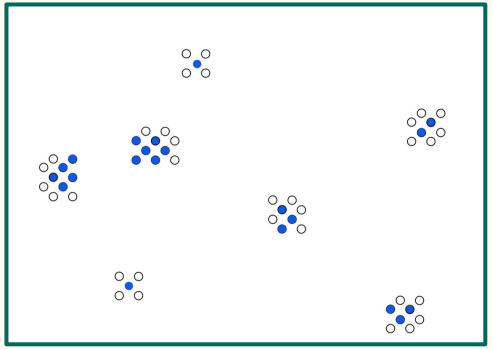

What About ML???

At a basic level, machine learning is about predicting the future based on the past. -Hal Daumé III

a set of methods that can automatically detect patterns in data, and then use the uncovered patterns to predict future data, or to perform other kinds of decision making under uncertainty - Kevin Murphy

Big Idea: Use Entropy/Information Gain to make “decisions”

Decision Trees

- Way to classify **importance** of a **feature** to the **label**
  - **Importance** - Use Entropy/Information Gain/Gain Ratio
  - **Feature** - characteristic of subject -> age, height, weight, etc.
  - **Label** - characteristic we are predicting -> disease, customer, etc.

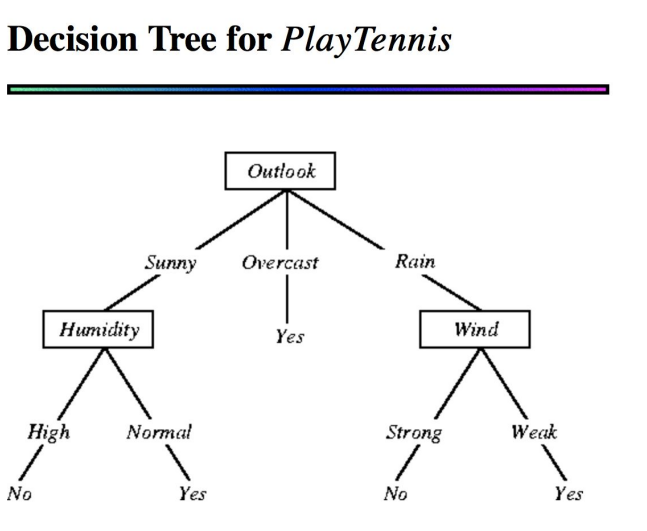

Information Gain/Gain Ratio/ Entropy

Quick Definitions:

Entropy -> Homogeneity, Information in System

Info Gain -> Reduction in Entropy from added Feature

Concept: pick features with highest info gain!

We want to remove uncertainty from our guess!!

Conditional Entropy

Specific Conditional Entropy:

A way to find the label given a single feature:

$$H(Y|X=v) = -\sum_{i \in vals(y)} P(Y=i|X=v) \log P(Y=i|X=v)$$

Overall:

$$H(Y|X) = \sum_{v \in vals(x)} P(X=v) H(Y|X=v)$$

Information Gain/Gain Ratio:

ID3 DTrees: Information Gain

- Make “decisions” using Info Gain
- Reducing Entropy - How much entropy did the feature remove from system
- For feature X, calculate:

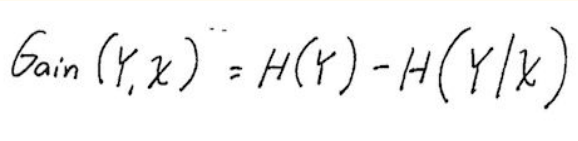

Highest Info Gain Used! Stop when all have same label .... But biased towards features with larger values

C4.5 DTrees: Gain Ratio

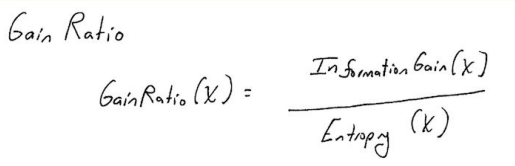

Result:

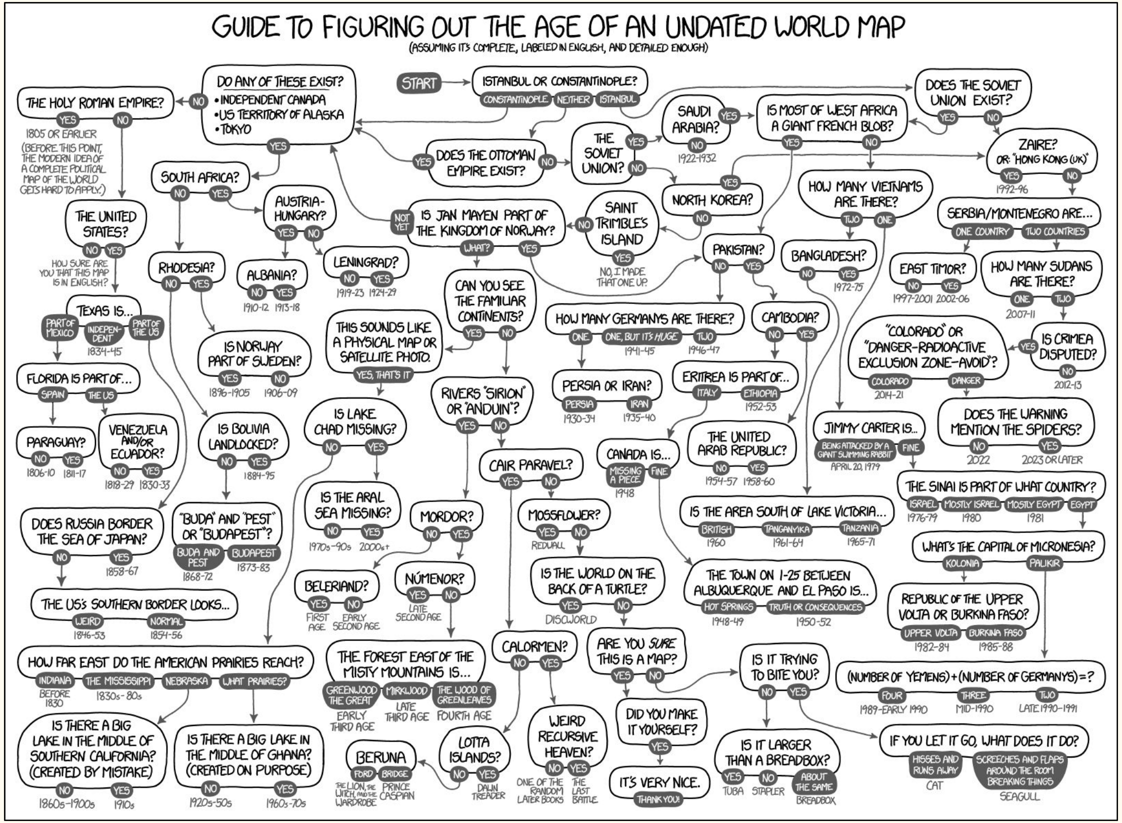

GUIDE TO FIGURING OUT THE AGE OF AN UNDATED WORLD MAP
(ASSUMING IT'S COMPLETE, LABELED IN ENGLISH, AND DETAILED ENOUGH)
START
ISTANBUL OR CONSTANTINOPLE?
CONSTANTINOPLE
NEITHER
ISTANBUL
DO ANY OF THESE EXIST?
• INDEPENDENT CANADA
• US TERRITORY OF ALASKA
• TOKYO
THE HOLY ROMAN EMPIRE?
1805 OR EARLIER
(BEFORE THIS POINT, THE MODERN IDEA OF A COMPLETE POLITICAL MAP OF THE WORLD GETS HARD TO APPLY.)
THE UNITED STATES?
HOW SURE ARE YOU THAT THIS MAP IS IN ENGLISH?
TEXAS IS...
PART OF MEXICO
INDEPENDENT
PART OF THE US
1834-45
FLORIDA IS PART OF...
SPAIN
THE US
PARAGUAY?
1806-10
1811-17
VENEZUELA AND/OR ECUADOR?
1818-29
1830-33
DOES RUSSIA BORDER THE SEA OF JAPAN?
1858-67
THE US'S SOUTHERN BORDER LOOKS...
WEIRD
NORMAL
1846-53
1854-56
HOW FAR EAST DO THE AMERICAN PRAIRIES REACH?
INDIANA
THE MISSISSIPPI
NEBRASKA
WHAT PRAIRIES?
BEFORE 1830
1830s-80s
IS THERE A BIG LAKE IN THE MIDDLE OF SOUTHERN CALIFORNIA? (CREATED BY MISTAKE)
1860s-1900s
1910s
IS THERE A BIG LAKE IN THE MIDDLE OF GHANA? (CREATED ON PURPOSE)
1920s-50s
1960s-70s
SOUTH AFRICA?
RHODESIA?
IS NORWAY PART OF SWEDEN?
1896-1905
1906-09
IS BOLIVIA LANDLOCKED?
1884-95
"BUDA" AND "PEST" OR "BUDAPEST"?
BUDA AND PEST
BUDAPEST
1868-72
1873-83
AUSTRIA-HUNGARY?
ALBANIA?
1910-12
1913-18
LENINGRAD?
1919-23
1924-29
DOES THE OTTOMAN EMPIRE EXIST?
IS JAN MAYEN PART OF THE KINGDOM OF NORWAY?
NOT YET
WHAT?
CAN YOU SEE THE FAMILIAR CONTINENTS?
THIS SOUNDS LIKE A PHYSICAL MAP OR SATELLITE PHOTO.
YES, THAT'S IT
IS LAKE CHAD MISSING?
IS THE ARAL SEA MISSING?
1970s-90s
2000s+
RIVERS "SIRION" OR "ANDUIN"?
MORDOR?
BELERIAND?
FIRST AGE
EARLY SECOND AGE
NÚMENOR?
LATE SECOND AGE
THE FOREST EAST OF THE MISTY MOUNTAINS IS...
GREENWOOD THE GREAT
MIRKWOOD
THE WOOD OF GREENLEAVES
EARLY THIRD AGE
LATE THIRD AGE
FOURTH AGE
CAIR PARAVEL?
MOSSFLOWER?
REDWALL
IS THE WORLD ON THE BACK OF A TURTLE?
DISCWORLD
CALORMEN?
BERUNA
FORD
BRIDGE
THE LION, THE WITCH, AND THE WARDROBE
PRINCE CASPIAN
LOTTA ISLANDS?
DAWN TREADER
WEIRD RECURSIVE HEAVEN?
ONE OF THE RANDOM LATER BOOKS
THE LAST BATTLE
ARE YOU SURE THIS IS A MAP?
DID YOU MAKE IT YOURSELF?
IT'S VERY NICE.
THANK YOU!
THE SOVIET UNION?
SAUDI ARABIA?
1922-1932
NORTH KOREA?
SAINT TRIMBLE'S ISLAND
NO, I MADE THAT ONE UP.
IS MOST OF WEST AFRICA A GIANT FRENCH BLOB?
DOES THE SOVIET UNION EXIST?
ZAIRE? OR: "HONG KONG (UK)"
1992-96
HOW MANY VIETNAMS ARE THERE?
TWO
ONE
PAKISTAN?
BANGLADESH?
1972-75
HOW MANY GERMANYS ARE THERE?
ONE
ONE, BUT IT'S HUGE
TWO
1941-45
1946-47
PERSIA OR IRAN?
PERSIA
IRAN
1930-34
1935-40
CAMBODIA?
ERITREA IS PART OF...
ITALY
ETHIOPIA
1952-53
CANADA IS...
MISSING A PIECE
FINE
1948
THE UNITED ARAB REPUBLIC?
1954-57
1958-60
IS THE AREA SOUTH OF LAKE VICTORIA...
BRITISH
TANGANYIKA
TANZANIA
1960
1961-64
1965-71
THE TOWN ON I-25 BETWEEN ALBUQUERQUE AND EL PASO IS...
HOT SPRINGS
TRUTH OR CONSEQUENCES
1948-49
1950-52
IS IT TRYING TO BITE YOU?
IS IT LARGER THAN A BREADBOX?
TUBA
STAPLER
ABOUT THE SAME
BREADBOX
JIMMY CARTER IS...
BEING ATTACKED BY A GIANT SWIMMING RABBIT
FINE
APRIL 20, 1979
SERBIA/MONTENEGRO ARE...
ONE COUNTRY
TWO COUNTRIES
EAST TIMOR?
1997-2001
2002-06
HOW MANY SUDANS ARE THERE?
2007-11
IS CRIMEA DISPUTED?
2012-13
"COLORADO" OR "DANGER-RADIOACTIVE EXCLUSION ZONE-AVOID"?
COLORADO
DANGER
2014-21
DOES THE WARNING MENTION THE SPIDERS?
2022
2023 OR LATER
THE SINAI IS PART OF WHAT COUNTRY?
ISRAEL
MOSTLY ISRAEL
MOSTLY EGYPT
EGYPT
1976-79
1980
1981
WHAT'S THE CAPITAL OF MICRONESIA?
KOLONIA
PALIKIR
REPUBLIC OF THE UPPER VOLTA OR BURKINA FASO?
UPPER VOLTA
BURKINA FASO
1982-84
1985-88
(NUMBER OF YEMENS) + (NUMBER OF GERMANYS) = ?
FOUR
THREE
TWO
1989-EARLY 1990
MID-1990
LATE 1990-1991
IF YOU LET IT GO, WHAT DOES IT DO?
HISSES AND RUNS AWAY
SCREECHES AND FLAPS AROUND THE ROOM BREAKING THINGS
CAT
SEAGULL

References

John McCarthy, Stanford University, http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html

Mackworth and Poole, Artificial Intelligence: Foundational Computational Agents

Swarthmore College Department of Computer Science

- CS63: Artificial Intelligence, Spring 2018, taught by Professor Wiedenbeck
- CS66: Machine Learning, Fall 2017, taught by Professor Soni

XKCD Comics