

AI, Machine Learning, Neural Networks and all that

Kevin Scannell
Saint Louis University
March 18, 2019

What is Artificial Intelligence (AI)?

- Simulation of human intelligence by computers
- Practitioners usually focus on particular tasks
- The Turing Test
- AI as a “moving target”

Famous Examples of AI tasks

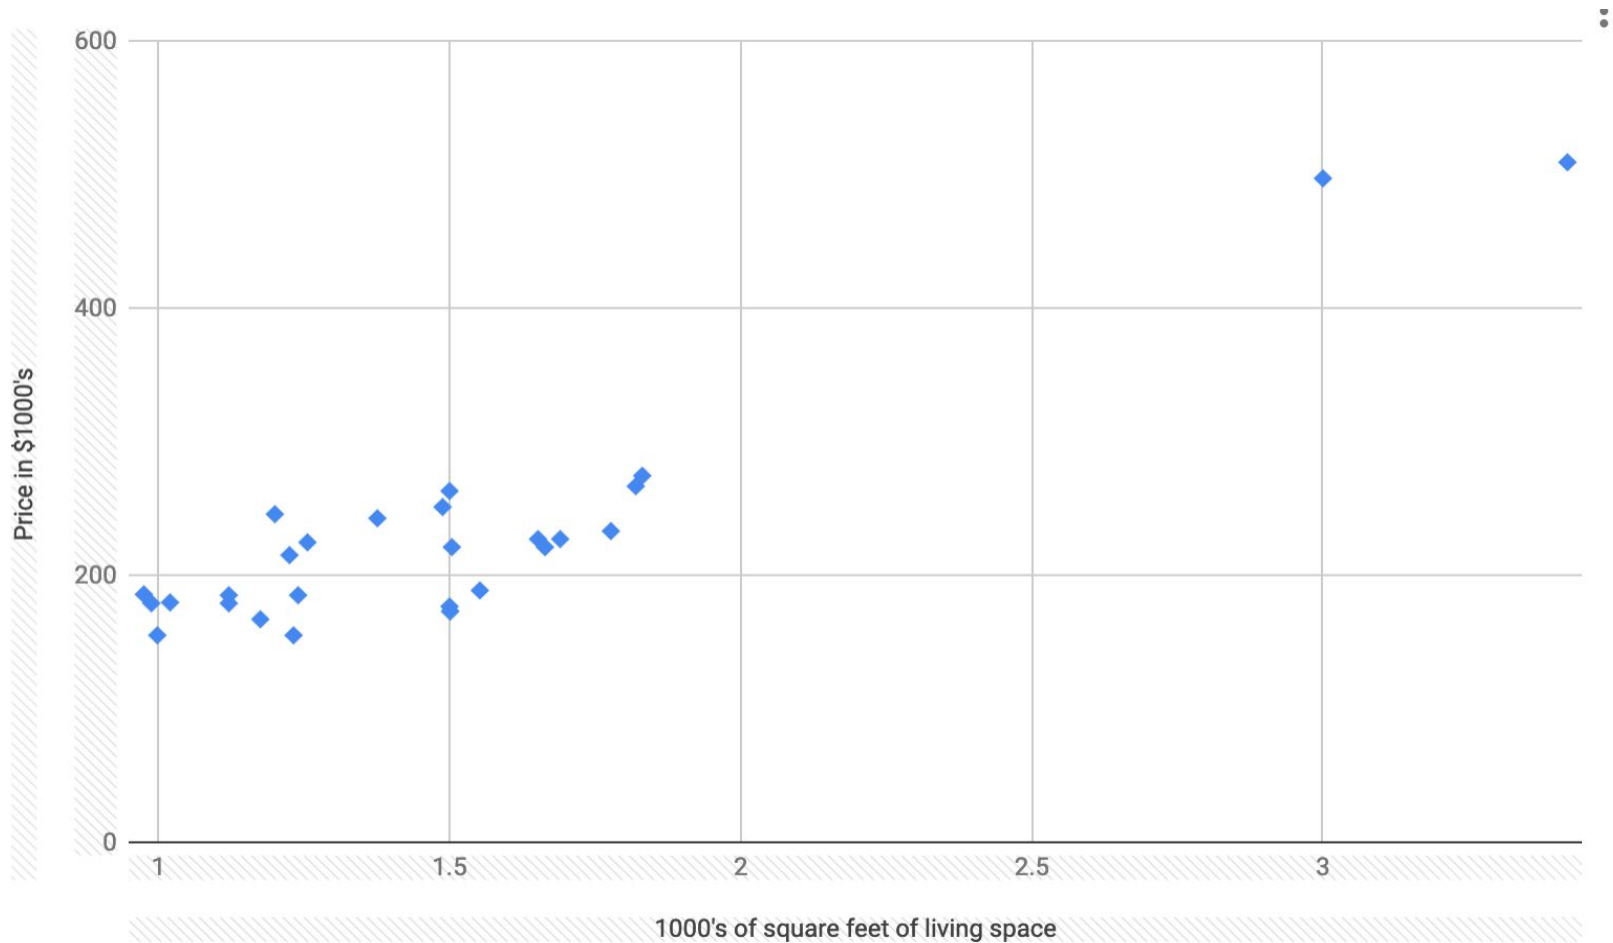
- Image processing (e.g. satellite imagery)
- Self-driving cars
- Computer Chess, Checkers, Go, etc.
- Question-answering (IBM Watson)
- Conversational agents, chatbots (Turing test)
- Multilingual interpreters, machine translation
- Stock market prediction

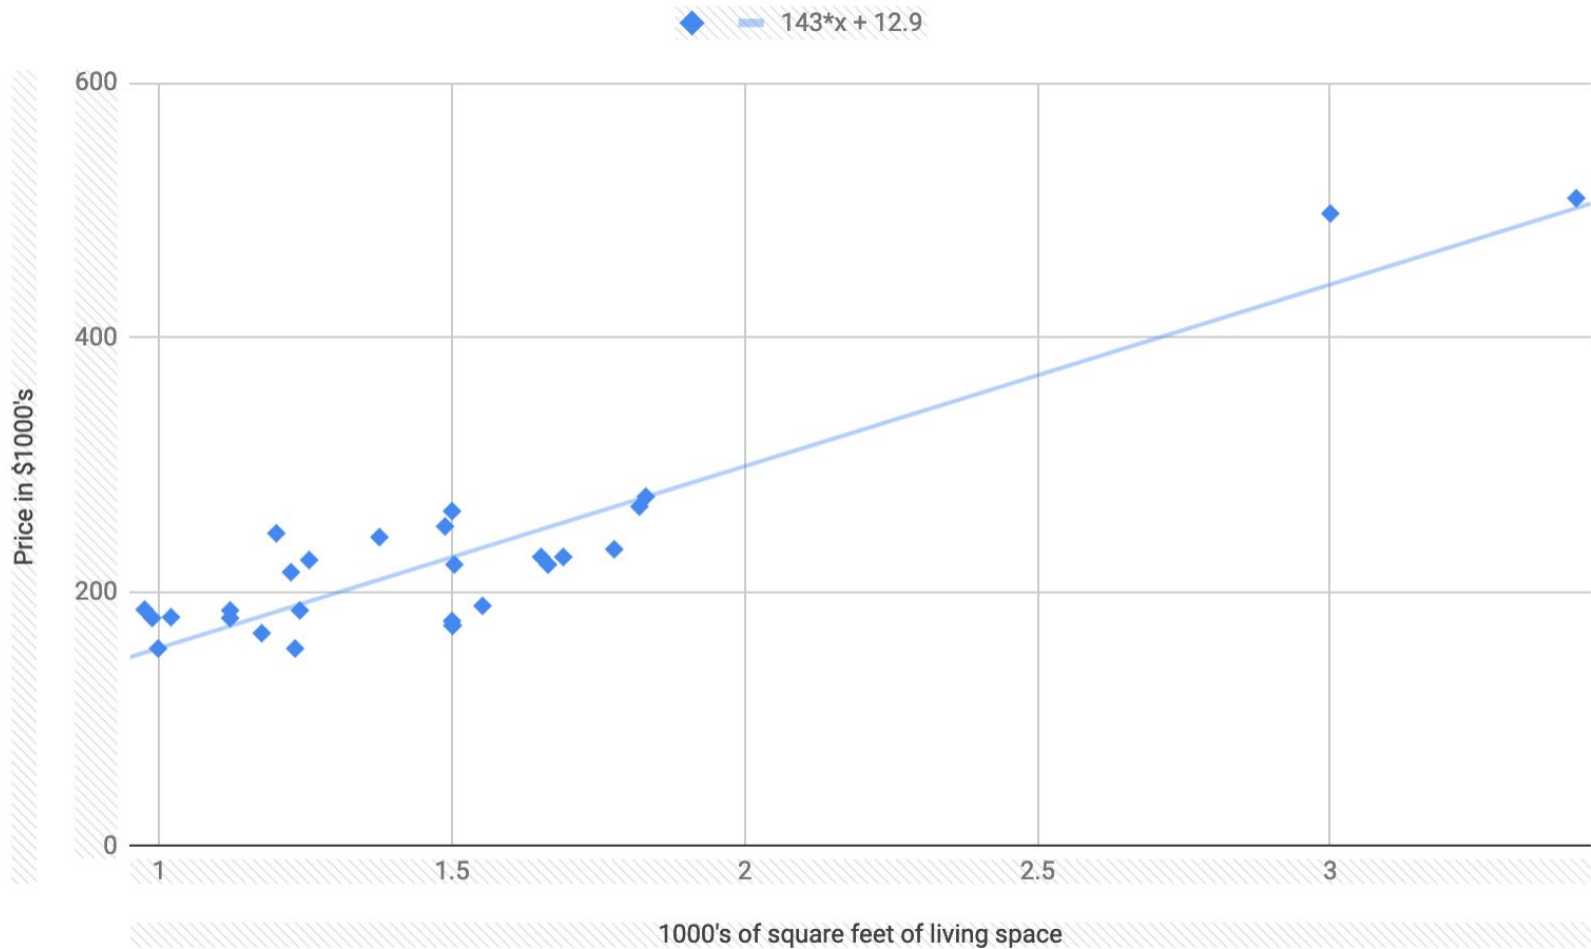
What is Machine Learning (ML)?

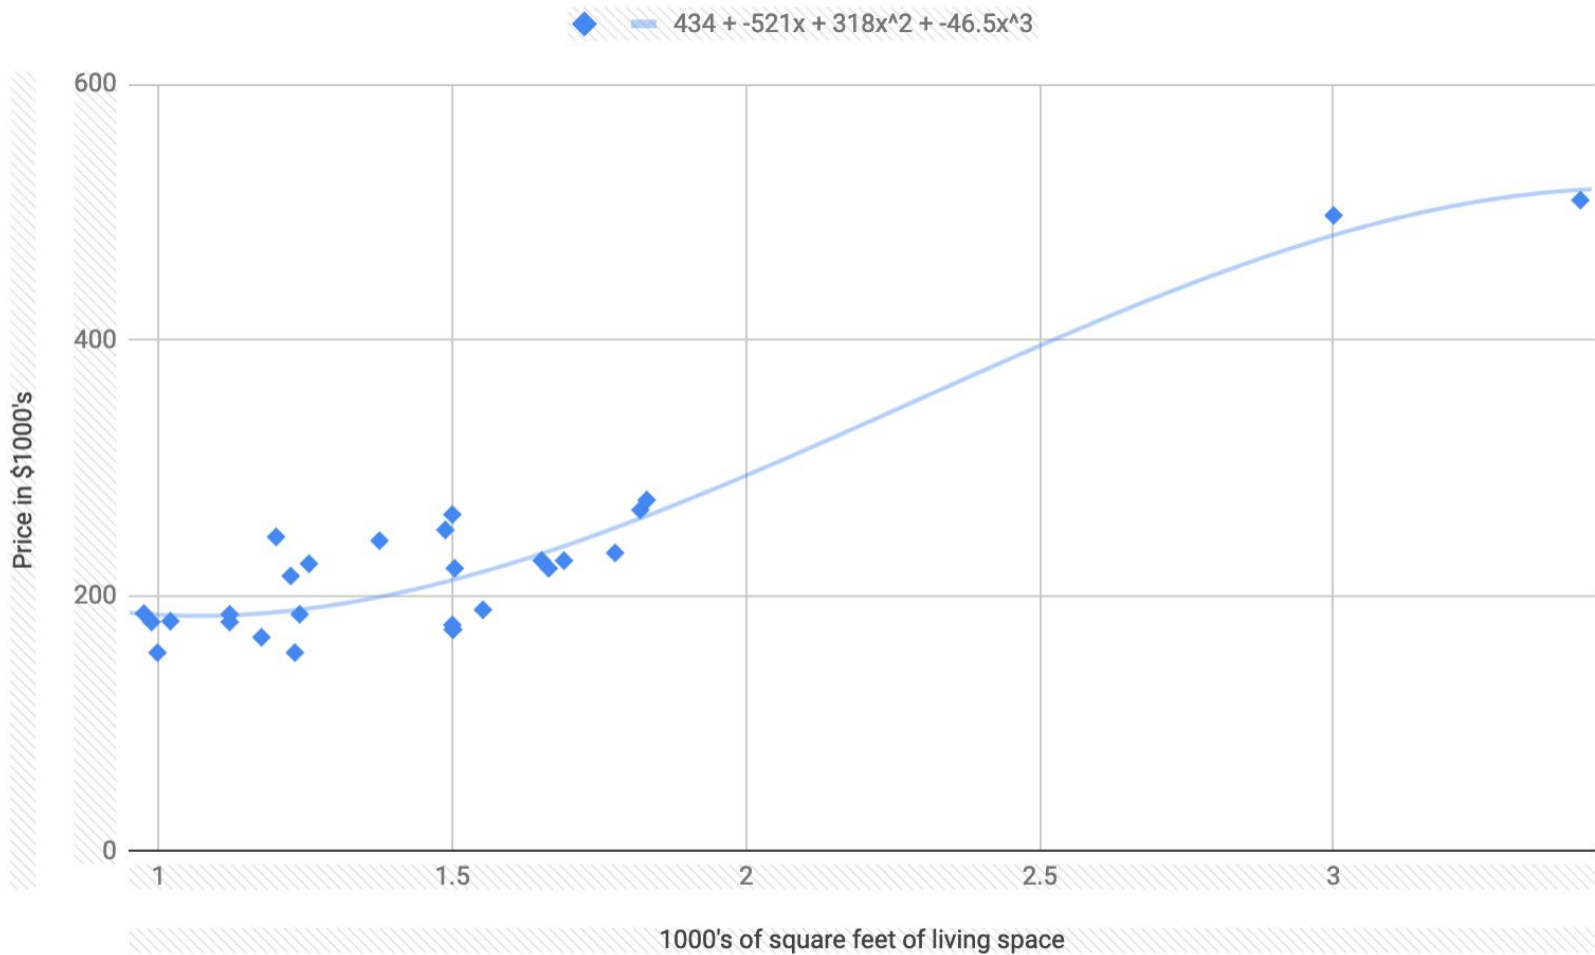
- One approach (the dominant one) to building AI systems
- Older systems often used rules programmed by humans
- With ML, the computer *learns* to perform a task
 - French > English machine translation for example
 - Computer is “shown” many sentences + translations
 - No need to know either language, or program grammar
 - Key: learn to *generalize* to never-seen examples

How much should you sell your house for?

- Gather a dataset of house prices
- Choose a model
- “Fit” that model to the data
- Make predictions based on the model
- Evaluate how good your model is










◆ $32837 - 122397x + 184225x^2 - 142147x^3 + 59140x^4 - 12544x^5 + 1060x^6$



What are neural networks?

- One approach (the dominant one) to machine learning
- Idea goes back to the dawn of computing, 1940's/50's
- Vast generalization of the example we just looked at
- They are models in this sense, millions/billions of params
- They are “fit” to data in exactly the same way
- Subject to overfitting/underfitting in the same way

What is a Language Model (LM)?

- Predicts probability of a word given the preceding words
- “I live in the United _____”
- “The glasses _____”
- “The man with the top hat and the glasses _____”
- “Please come to my party on _____”
- A (not very good) LM is probably built into your phone!

Using language models generatively

- Choose next word according to probabilities from the LM
- Can condition the probabilities on a prompt, a source language text, or a conversational turn
- Better LMs => better chatbots, machine translation, etc.

OpenAI's GPT-2

- OpenAI is a non-profit AI research organization
- Founded by Elon Musk and (St. Louisan) Sam Altman
- One month ago announced a new LM for English
- Trained on roughly *7 billion* words from the web
- Significant leap forward by several objective measures
- Qualitative evaluation by using it generatively...

SYSTEM PROMPT
(HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

MODEL
COMPLETION
(MACHINE-
WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

Ethics and AI

- OpenAI has decided *not* to release the model or datasets
- “...due to concerns about [the model] being used to generate deceptive, biased, or abusive language at scale”
- Have promised to revisit the issue in six months