CSC 580: Principles of Machine Learning

Spring 2020

Why Machine Learning? (AI?)

“Artificial intelligence is the science of making machines do things that would require intelligence if done by men”
Marvin Minsky, … 1968


1997
IBM
RS6000 SP

ARTICLE

doi:10.1038/nature24270

Mastering the game of Go without human knowledge

David Silver[1]*, Julian Schrittwieser[1]*, Karen Simonyan[1]*, Ioannis Antonoglou[1], Aja Huang[1], Arthur Guez[1], Thomas Hubert[1], Lucas Baker[1], Matthew Lai[1], Adrian Bolton[1], Yutian Chen[1], Timothy Lillicrap[1], Fan Hui[1], Laurent Sifre[1], George van den Driessche[1], Thore Graepel[1] & Demis Hassabis[1]

A long-standing goal of artificial intelligence is an algorithm that learns, *tabula rasa*, superhuman proficiency in challenging domains. Recently, AlphaGo became the first program to defeat a world champion in the game of Go. The tree search in AlphaGo evaluated positions and selected moves using deep neural networks. These neural networks were trained by supervised learning from human expert moves, and by reinforcement learning from self-play. Here we introduce an algorithm based solely on reinforcement learning, without human data, guidance or domain knowledge beyond game rules. AlphaGo becomes its own teacher: a neural network is trained to predict AlphaGo's own move selections and also the winner of AlphaGo's games. This neural network improves the strength of the tree search, resulting in higher quality move selection and stronger self-play in the next iteration. Starting *tabula rasa*, our new program AlphaGo Zero achieved superhuman performance, winning 100–0 against the previously published, champion-defeating AlphaGo.

Much progress towards artificial intelligence has been made using supervised learning systems that are trained to replicate the decisions of human experts[1–4]. However, expert data sets are often expensive, unreliable or simply unavailable. Even when reliable data sets are available, they may impose a ceiling on the performance of systems trained in this manner[5]. By contrast, reinforcement learning systems are trained from their own experience, in principle allowing them to exceed human capabilities, and to operate in domains where human expertise is lacking. Recently, there has been rapid progress towards this goal, using deep neural networks trained by reinforcement learning. These systems have outperformed humans in computer games, such as Atari[6,7] and 3D virtual environments[8–10]. However, the most chal-

trained solely by self-play reinforcement learning, starting from random play, without any supervision or use of human data. Second, it uses only the black and white stones from the board as input features. Third, it uses a single neural network, rather than separate policy and value networks. Finally, it uses a simpler tree search that relies upon this single neural network to evaluate positions and sample moves, without performing any Monte Carlo rollouts. To achieve these results, we introduce a new reinforcement learning algorithm that incorporates lookahead search inside the training loop, resulting in rapid improvement and precise and stable learning. Further technical differences in the search algorithm, training procedure and network architecture are described in Methods.

Spam Filtering

Precision Toyota of.	**The 2019 Prius packs an all-wheel wallop.** - Discover the many features of the 2019 Prius. View With Ima
Amy Green	**Western blot automation: free trial** - BlotCycler www.blotcycler.com Blot-to-blot consistency. Time after t
Christopher	**Can you do me a favor?** - We're nominated for Best Airport Service / Amenity at Dallas Fort Worth Can yo
Christopher	**USA Today Readers' Choice 2019** - We're nominated for Best Airport Service / Amenity at Dallas Fort Wor
Meilleurtaux	**Jusqu'à -60% sur vos mensualités** - _ Regrou
Enzo Life Sciences	**A New Notebook for Your New Year's Resolutions!** - Update your email preferences. Decide what we senc
One Budget	**Simulez votre rachat de crédit en 3minutes** - _
Comme j'aime	**En 2019, Je prends de bonnes résolutions** - _
Brian Nosek	**Last chance to support COS, OSF, and open science in 2018!** - View this email in your browser Dear Carlo
M. M. Fridman	**Free Will Offer** - -- I, Mikhail Fridman have selected you specifically as one of my beneficiaries for my Cha
Enzo Life Sciences	**Wishing You a Joyous Holiday Season** - Update your email preferences. Decide what we send & how ofte
Carlos	**Hello! remember me? I'm Carlos (ex coworker)** - Hello! dear friend I am Carlos, your ex. co-worker remem
Morwen R	**Hi, remeber me? I am Morwen, ex.coworker** - Hello ! I'm Morwen, your ex. co-worker Do you remeber me?
Assurance Animaux	**Offre spécial web: 3 mois offerts à partir de 8,60 par mois**** - _
LinkedIn	**You appeared in 9 search this week** - You appeared in 9 search this week LinkedIn carlosscheidegger@gr
infovis-request	**Infovis Digest, Vol 151, Issue 7** - Send Infovis mailing list submissions to infovis@infovis.org To subscrib
infovis-request	**Infovis Digest, Vol 151, Issue 6** - Send Infovis mailing list submissions to infovis@infovis.org To subscrib
Southwest Gas	**We've Updated Our Terms & Conditions and Privacy Policy** - Southwest Gas UPDATES for you OUR Terms
infovis-request	**Infovis Digest, Vol 151, Issue 5** - Send Infovis mailing list submissions to infovis@infovis.org To subscrib
Christopher at Minu.	**We're making headlines again!** - 'Tis the season for merry-making, and we have some exciting news to sh

What do we have?

- Fast computers! We could simply write the whole spam-filtering programs ourselves
  - But (for example) spam changes over time, and we wish we didn't have to rewrite our programs every day
  - And it's really hard to write code that works well
- **Given any program, we can say how well a program does on data that is known good (our "training set")**

What’s our plan?

- We make the computer write its own programs

- More strictly, we write *meta-programs* that are easy for a computer to optimize
    - It’s easier for computers to search over regular structure than irregular, so the space of programs computers search over is “simpler” than the space of programs humans search over.

What’s our plan?

- In ML, we study how to organize software so that it’s easy to find a good program to use for our task
- Different definitions of “easy”, “find”, “best”, and “task” account for most of what we will see in the course

What’s our problem?

- Are we trying to predict (“regress on”) tomorrow’s temperature in Tucson?
- Are we trying to label (“classify”) an image?
    - Is there more than one label?
- Are we trying to generate an image from a label?
- Are we trying to translate a sentence?
- Are we trying to decide whether a bank customer gets a loan?

What are our tools?

- Probability
- Statistics
- Calculus
- Linear Algebra
- Optimization
- Software Engineering

Structure of the Course

- Weekly readings and homework assignments
  - We'll use Hal Daume's CIML as the main text (http://ciml.info)
  - Some readings from Bishop's PRML (https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf - just google it)
- Final project
- Midterm, Final Exam

- This is an intense course, be ready.

PRML, 1.1, 1.2

Next Lecture

- Required reading:
  - CIML, chapter 1 (Decision Trees)
  - PRML, 1.1, 1.2.4, 1.2.5, 1.2.6