Actor Critic Methods: From Paper to Code

What You Will Learn In This Course

A New Path

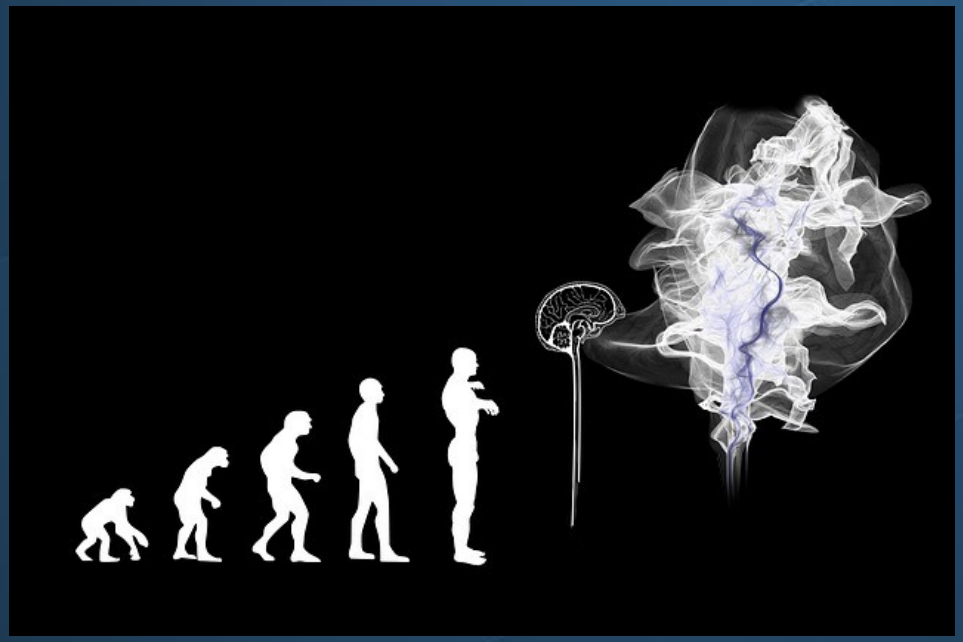

Route to AGI?

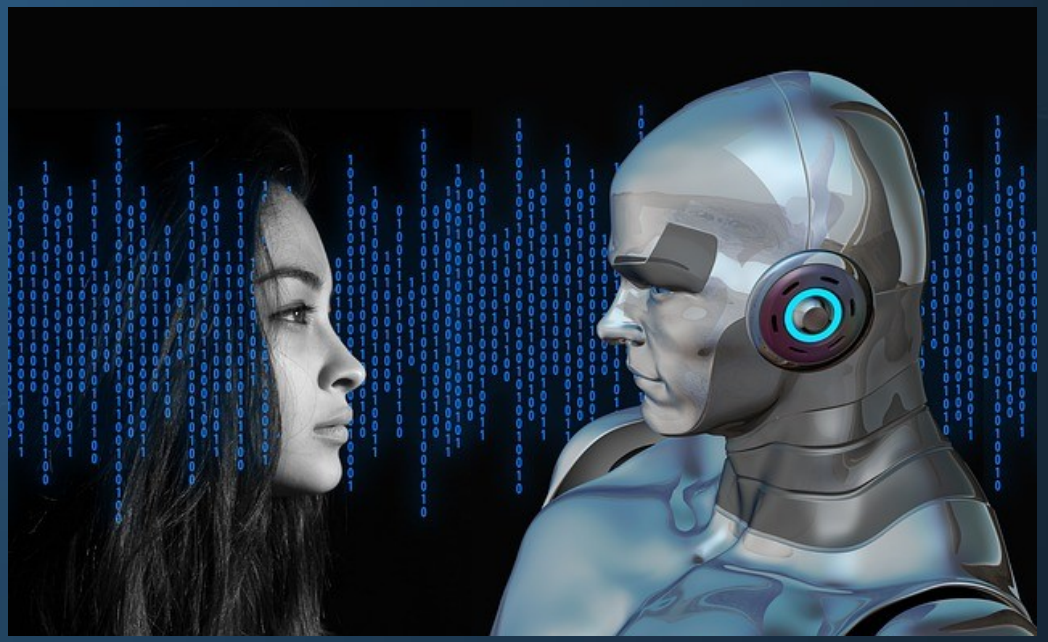

Mimic human learning

Not Just a Toy

Deep Reinforcement Learning for Join Order Enumeration

Ryan Marcus
Brandeis University
ryan@cs.brandeis.edu

Olga Papaemmanouil
Brandeis University
olga@cs.brandeis.edu

ABSTRACT

Join order selection plays a significant role in query performance. However, modern query optimizers typically employ static join order enumeration algorithms that do not incorporate feedback about the quality of the resulting plan. Hence, optimizers often repeatedly choose the same bad plan, as they have no mechanism for "learning from their mistakes." Here, we argue that deep reinforcement learning techniques can be applied to address this challenge. These techniques, powered by artificial neural networks, can *automatically* improve optimizer decision-making by incorporating feedback. Towards this goal, we present *ReJOIN*, a proof-of-concept join enumerator, as well as preliminary results indicating that ReJOIN can match or outperform the PostgreSQL optimizer in terms

Unfortunately, these heuristics often miss good execution plans. More importantly, traditional query optimizers rely on *static* strategies, and hence do not learn from past experiences. Traditional systems plan a query, execute the query plan, and then forget they ever optimized this query. Because of the lack of feedback, a query optimizer may select the same bad plan repeatedly, never learning from its previous good or bad choices.

In this paper, we share our vision of a *learning-based* optimizer that leverages prior experience, aiming to learn how to optimize fu-

Active Object Localization with Deep Reinforcement Learning

Juan C. Caicedo
Fundación Universitaria Konrad Lorenz
Bogotá, Colombia
juanc.caicedor@konradlorenz.edu.co

Svetlana Lazebnik
University of Illinois at Urbana Champaign
Urbana, IL, USA
slazebni@illinois.edu

Abstract

We present an active detection model for localizing objects in scenes. The model is class-specific and allows an agent to focus attention on candidate regions for identifying the correct location of a target object. This agent learns to deform a bounding box using simple transformation actions, with the goal of determining the most specific location of target objects following top-down reasoning. The proposed localization agent is trained using deep reinforcement learning, and evaluated on the Pascal VOC 2007 dataset. We show that agents guided by the proposed model are able to localize a single instance of an object after analyzing only between 11 and 25 regions in an image,

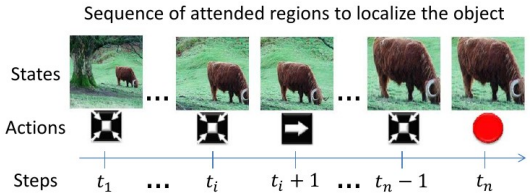

Figure 1. A sequence of actions taken by the proposed algorithm to localize a cow. The algorithm attends regions and decides how to transform the bounding box to progressively localize the object.

tures to achieve state-of-the-art results in the Pascal and ImageNet benchmarks. Several other works have proposed

What You Will Learn

Implement papers

Reinforcement learning from scratch

Reinforcement Learning

Theoretical fundamentals

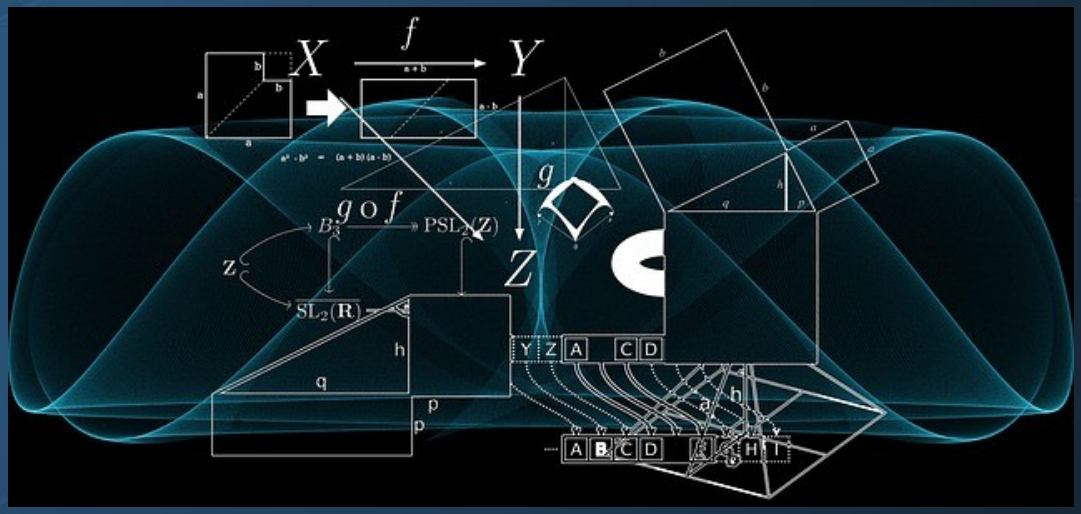

Applications in Open AI gym

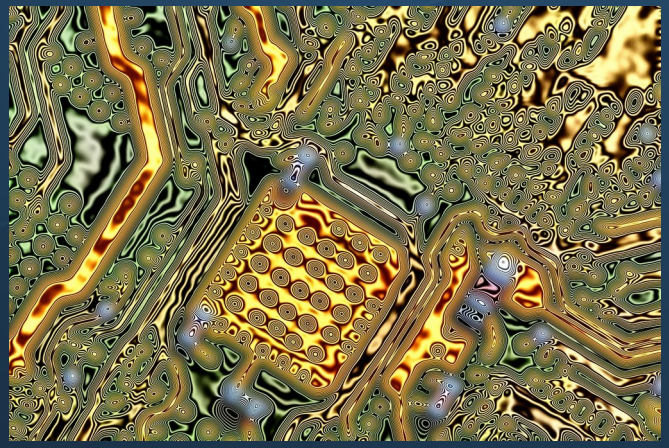

Deep reinforcement learning

Deep Deterministic Policy Gradients

Published as a conference paper at ICLR 2016

CONTINUOUS CONTROL WITH DEEP REINFORCEMENT LEARNING

Timothy P. Lillicrap,* Jonathan J. Hunt,* Alexander Pritzel, Nicolas Heess,
Tom Erez, Yuval Tassa, David Silver & Daan Wierstra
Google Deepmind
London, UK
{countzero, jjhunt, apritzel, heess,
etom, tassa, davidsilver, wierstra} @ google.com

ABSTRACT

We adapt the ideas underlying the success of Deep Q-Learning to the continuous action domain. We present an actor-critic, model-free algorithm based on the deterministic policy gradient that can operate over continuous action spaces. Using the same learning algorithm, network architecture and hyper-parameters, our algorithm robustly solves more than 20 simulated physics tasks, including classic problems such as cartpole swing-up, dexterous manipulation, legged locomotion

Focus on Critical Skills

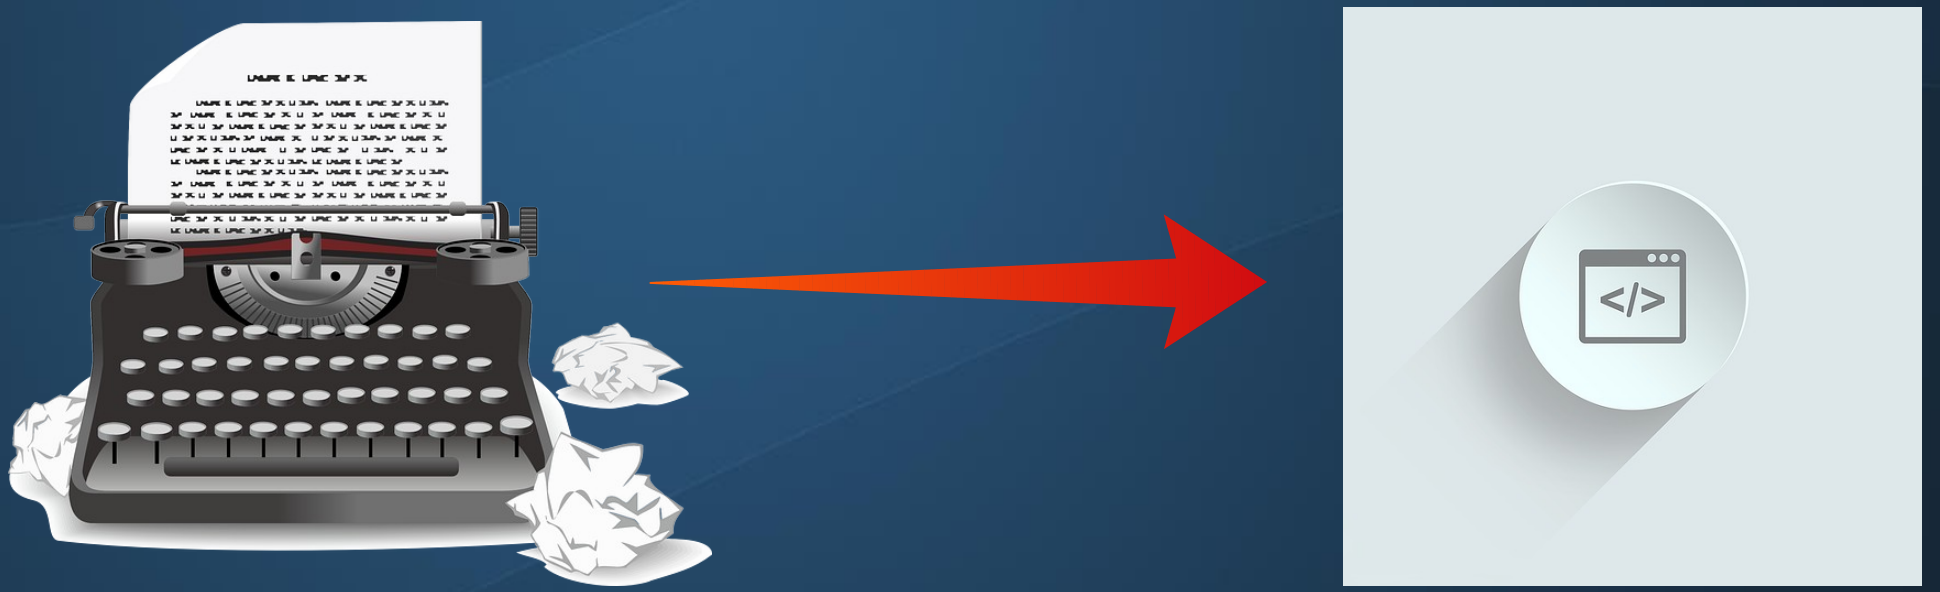

Paper to code

Focus on Fundamentals

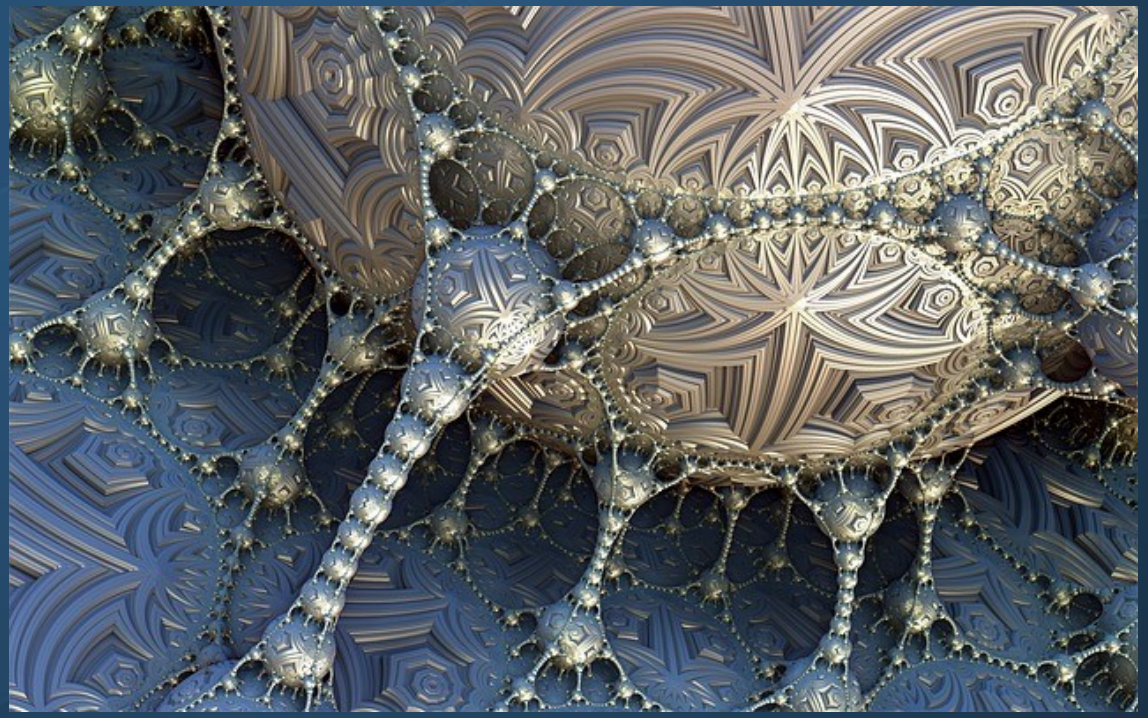

Mathematics is critical

Some Caveats

- Focus on demonstrating learning
- Due to huge amount of resources

Up Next