Machine Lifelong Learning: Challenges and Benefits for Artificial General Intelligence

Zach Cao, Alex Herzog, Ian Stedham

Intro

Background

- Machine learning has been steadily advancing for quite some time
- These algorithms get trained for one purpose
- There should be more focus on algorithms that can learn for their lifetime

Motivation

- New theories of learning are being developed that support this approach
  - These theories include constructive induction, sequential task learning, and learning with deep architectures
- Computational and storage power only continues to increase
- Challenging but beneficial field of study

Lifelong learning

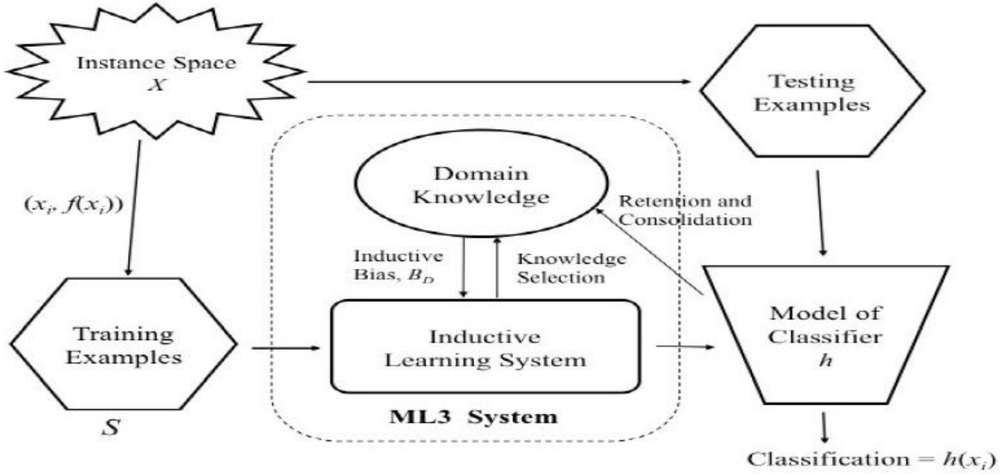

Lifelong Learning Framework

- Inductive bias is necessary for lifelong learning
- Used to develop new hypotheses from general examples

Domain Knowledge

- Initial training of lifelong learning models is identical to other inductive learning systems
- Knowledge from each hypothesis made is saved in Domain Knowledge
- Domain Knowledge is a long term memory structure
- Domain Knowledge is used to provide beneficial inductive bias to new tasks

Inductive bias

What is inductive bias

- The constraint on a learning system's hypothesis space, beyond the criterion of consistency with the training examples.
- Essential for generalization or accurate extrapolation

Using inductive bias

- Selecting aspects of domain knowledge can help guide inductive bias
  - Selecting the most relevant and helpful information is a challenge
- Built on the transfer of knowledge from one task to another
- Systems conducting their own search for relevant information is core to learning and strengthens hypothesis accuracy

Guiding inductive bias

Methods:

- Long and short term memory
  - Helps systems to overcome the stability-plasticity problem
- Constructive induction learning
  - Learning process that finds a representation space and definition for the concept of the space
- Incremental learning
  - Learning from small primitive tasks
- Systems with high explainability
- Task rehearsal
  - Maintains accuracy in previous tasks and helps knowledge transfer between older and newer tasks
- Continual learning
  - Uses hierarchical and incremental learning to build on previously learned tasks
- Deep neural networks

Benefits and Challenges

- **Pros**
    - Application of transfer learning
        - Prior training can improve learning new tasks
    - Unsupervised learning
    - Knowledge retention
    - Inductive Transfer
        - The system can prioritize knowledge
    - Task Generalization
- **Cons**
    - Relevance of knowledge
        - Potentially Subjective
    - Stability-Plasticity Problem
        - Learning new non-competing information
    - Scalability
        - Inputs and operations will scale exponentially
    - Labor Intense Research

Relevance to AGI

- Machine lifelong learning aims to model a similar way to how humans think
- Many of the challenges faced by MLL will be obstacles in the development of AGI
- MLL is one facet of how the human brain operates (which AGI will model)
  - Understanding MLL might provide some insight into the nature of AGI

Questions and Discussion

How do you accurately classify relevant information?

Could AI gain general intelligence from MLL systems?

To what degree can MLL systems overcome bias originating from data discrepancies?

Reference

https://link.springer.com/chapter/10.1007/978-3-642-22887-2_45