T-720-ATAI-2016

Advanced Topics in Artificial Intelligence:

Universal Pedagogy

Jordi Bieger

School of Computer Science | Center for Analysis and Design of Intelligent Agents

Pedagogy

- Pedagogy
  - The study of how to teach...
  - …humans.
- Artificial Pedagogy
  - The study of how to teach artificial learning systems.
- Universal Pedagogy
  - The study of how to teach learning systems in general.
  - Including humans, other animals and AI.

Key Questions

- Given information about a learner, teaching goal and constraints, what is the best way to teach?
- What teaching methods are there, and when are they applicable?
- How can we evaluate the learner and the teacher?

Comparison with ML

- No theory of teaching exists
- Focus on teaching rather than learning
- Aimed at full spectrum of learning systems
- Emphasis towards AGI-aspiring systems

Why Artificial Pedagogy?

- Teaching is helpful but complex
- Current machine teaching is ad hoc
- Sophisticated teaching needed in complex domains
- Sufficiently advanced learners now exist
- Relevance will increase as AI field advances

General Intelligence

- The ability to operate successfully in a wide range of complex environments not known at design time.
- Implicit:
  - Wide range of goals, tasks and problems
  - Constant learning and adaptation
  - Assumption of Insufficient Knowledge and Resources (AIKR)

Knowledge Acquisition

- Successful operation requires relevant knowledge
  - declarative (beliefs), procedural (skills), structural (priorities)
- Knowledge programming:
  - knowledge infusion by system designers before the system starts its operational lifecycle
- Learning:
  - the acquisition and adaptation of knowledge during the system's "life"
- Teaching:
  - actions aimed at helping another agent learn

Behavior

- Behavior is determined by both design and experience: nature and nurture.

- Most focus is on controllable parts of nature over nurture and structure over content.

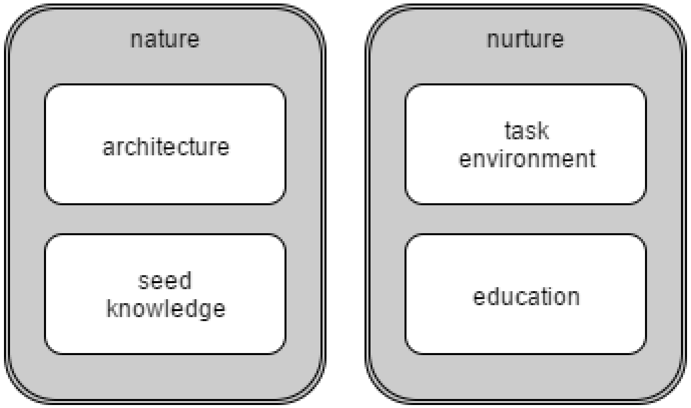

Teaching vs. Knowledge Programming

- Programming:
  - minimal seed knowledge required
  - precise

- Teaching:
  - natural
  - adaptive
  - on-the-fly
  - can't program everything

Learning Systems

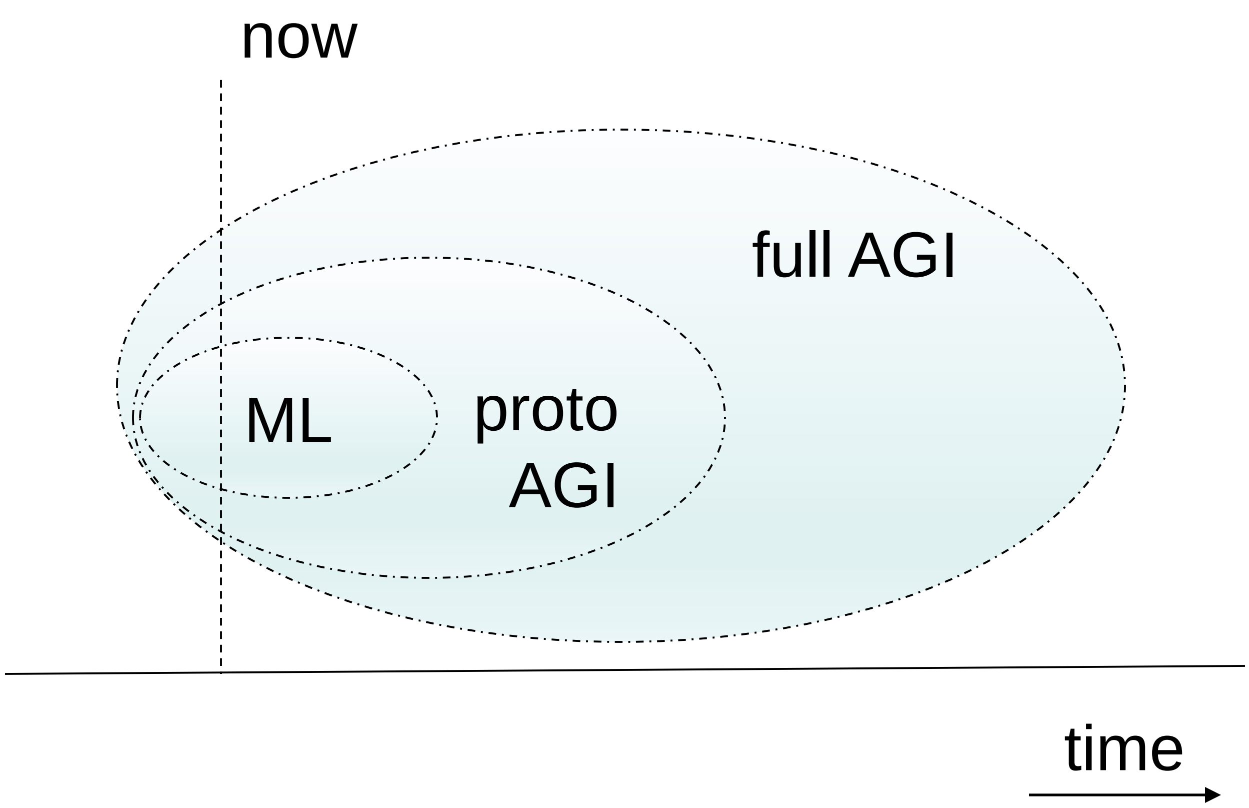

Learning Systems

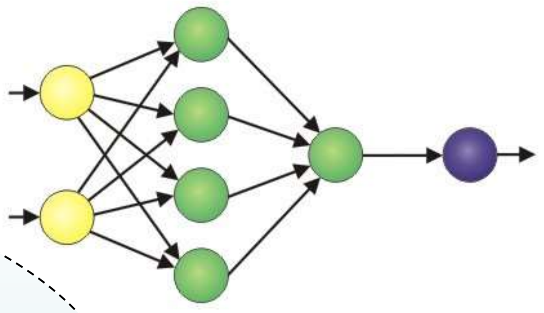

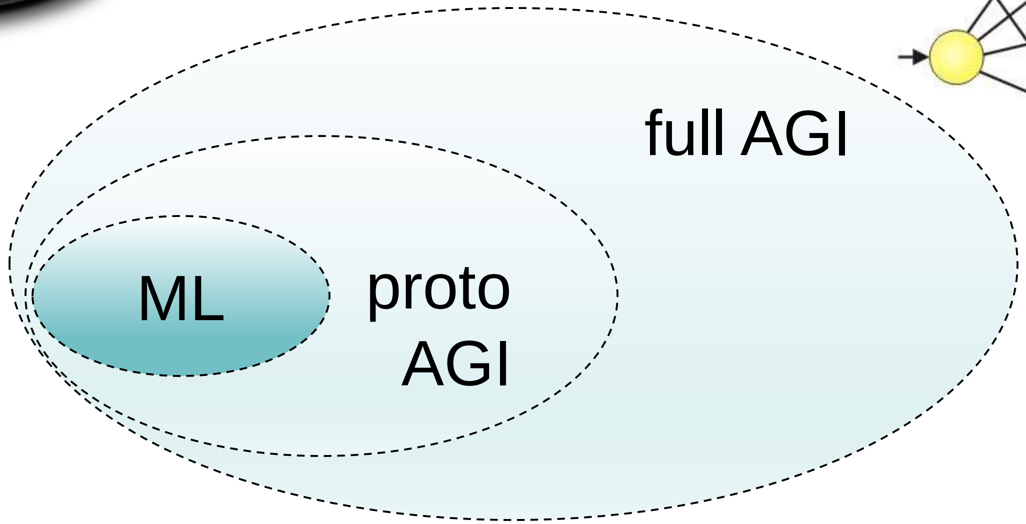

Systems:
ANNs
SVMs
Q-learning
K-means
…

Applicable in specialized simplified domains

Machine Learning

- Unsupervised Learning
  - problems: clustering, dimensionality reduction, visualization
  - teaching: typically none
- Reinforcement Learning
  - problems: optimal control
  - teaching: using reward and punishment
- Supervised Learning
  - problems: prediction, classification, regression
  - teaching: by example

AP vs. Supervised Learning

- Specific kind of teaching: by example.
- Input-target pairs are provided directly to the learner, using the learner's exact interface.
- Feedback is immediate, extremely detailed and applied to atomic decisions.
- Usually no online learning.
- Usually very simple, passive learners incapable of online, multitask or transfer learning.
- In reality communication can exist at different time scales and levels of abstraction and the learner's exact interfaces may be unknown or incompatible.

Example: Neural Network

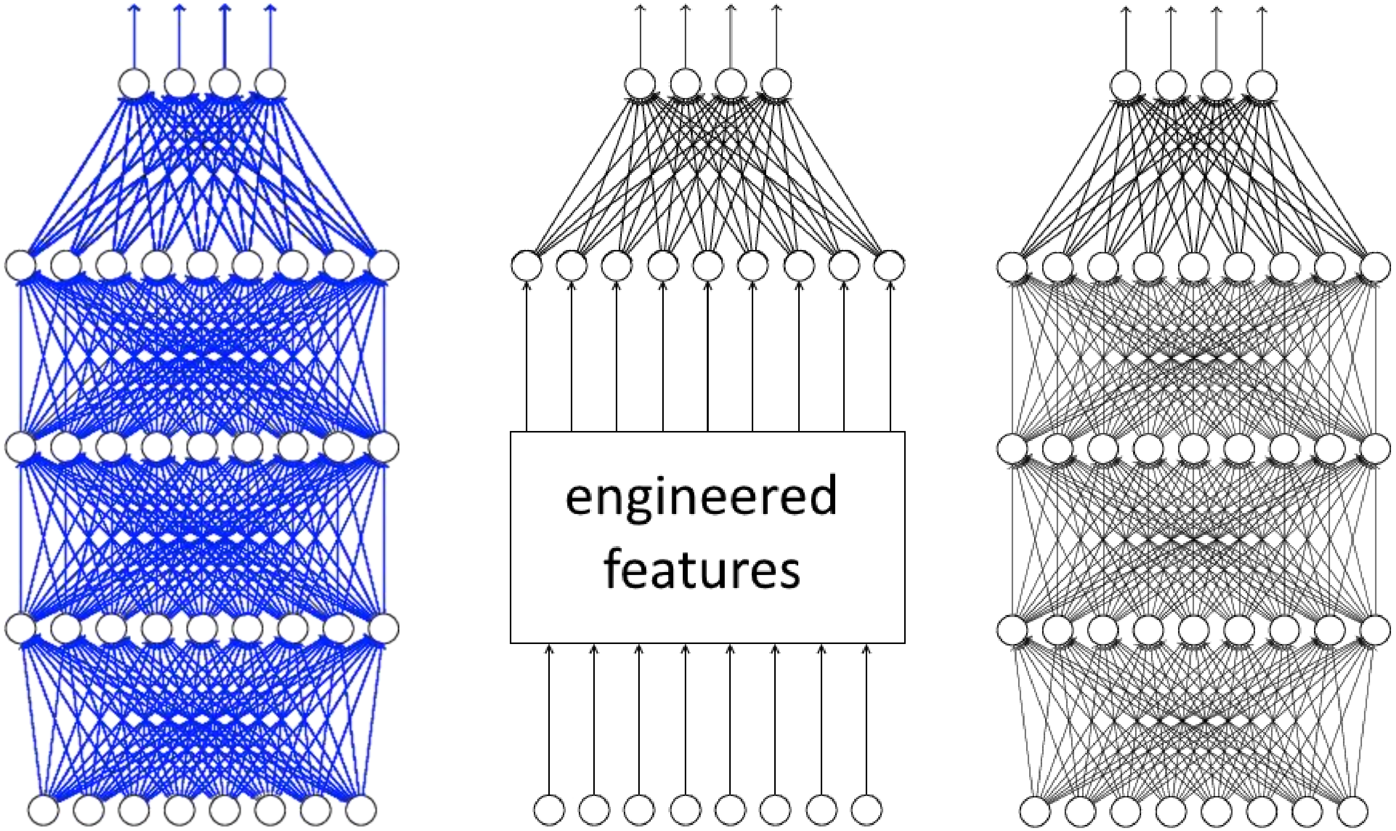

manual weight coding

feature engineering

deep learning

Machine Teaching

- Zhu (2015) recognized the importance of teaching , coined the phrase ‘machine teaching’:
  - “the problem of finding an optimal training set given a machine learning algorithm and a target model” (Zhu, 2015)

- Focused on traditional supervised learning
  - “generating fascinating mathematical questions for computer scientists to ponder” (Zhu, 2015)

  - c.f. Zhu, X. (2015). Machine Teaching: An Inverse Problem to Machine Learning and an Approach Toward Optimal Education. In *29th AAAI Conference*. AAAI.

Teaching in AI

- Curriculum learning / shaping
- Task decomposition
- Demonstration / imitation
- Teleoperation

- Still ad hoc
  - c.f. Bieger, J., Thórisson, K. R., & Garrett, D. (2014). Raising AI: Tutoring Matters. In *Artificial General Intelligence* (pp. 1-10). Springer.

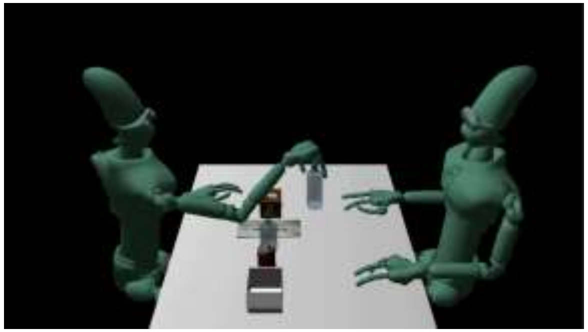

Learning Systems

Abilities:
lifelong learning
transfer learning
AIKR
real-time

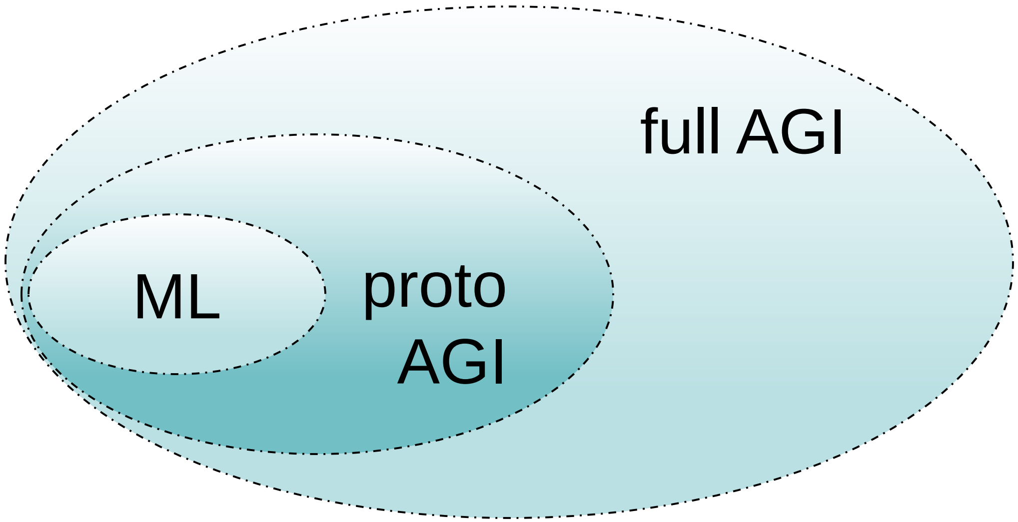

Systems:
AERA
NARS
iCub?
OpenCog?
...

Sophisticated Teaching Required

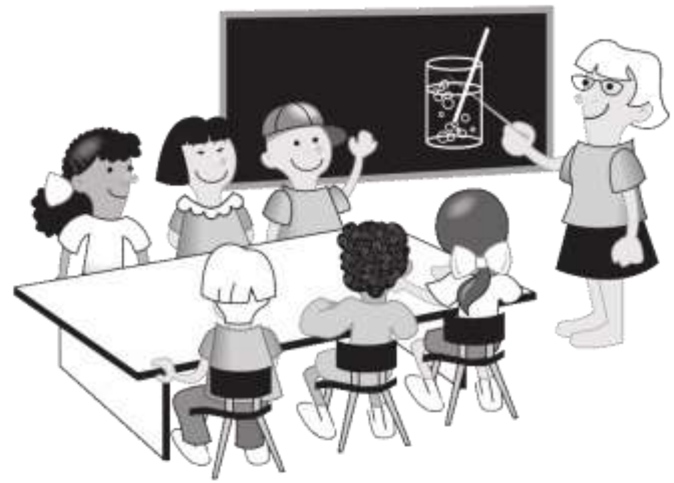

	reinforcement		verbal instruction	
	pull lever	read	pull lever	read
mouse	✓	✗	✗	✗
human	✓	✗	✓	✓

Developmental AI

- Goal: Cognitive ability grows through interaction
- Inspired by cognitive and developmental psychology
- Focus on learner
- Ignores pedagogy

Cognitive Architectures

- Increased intelligence calls for complex systems
- Diverse landscape:
  - AERA, NARS, LIDA, CogPrime, MicroPsi, …
- Increased impact of general theory

Learning Systems

Abilities:
autonomy
flexibility
analogies
abstraction
reflection
…

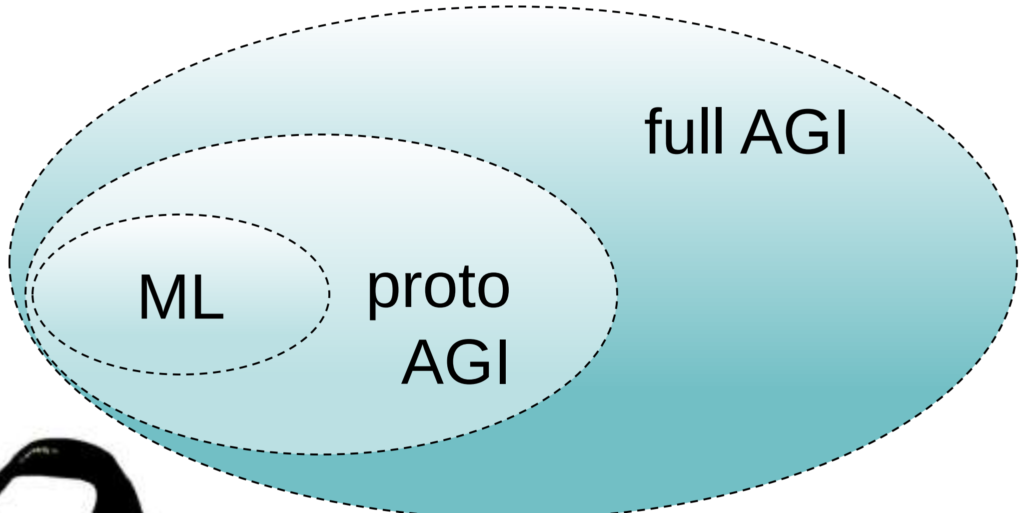

Systems:
Humans
Future AI?
…

Requirements on Learners

- Lifelong learning
  - incremental / online / continuous
  - multitask
- Transfer learning
- Active learning
- Real-time interaction

Key Questions

- Given information about a learner, teaching goal and constraints, what is the best way to teach?
- How can we evaluate the learner and the teacher?
- What teaching methods are there, and when are they applicable?

Modeling the Task-Environment

- We need a characterization of task-environments in terms of fundamental properties (e.g. complexity, observability, controllability, …).
- This would allow for the comparison of task-environments and learners' performance.
- We need some way to build environments with particular properties.
- Ideally the teacher should be able to create task-environments that allow for meaningful interactive evaluation and that help the learner improve.

Modelling the Learner

- Teaching requires having a model of the learning audience.
- It is unrealistic to have full knowledge of the learner's implementation, current knowledge and performance.
- It is undesirable to base all decisions on unfounded assumptions.
- Interactive evaluation is needed.
- We would like to characterize the learner in terms of fundamental properties and capabilities (e.g. reliability, resistance, learning speed and capacity, current performance, …)

Evaluation

- Turing test, Algorithmic IQ, General Game Playing...
- Measure single point in time.
- Cannot measure growth.
- Freeze time.
- Specific, unrepresentative and unmodifiable tasks.
- Procedurally generated tasks are too abstract, and don't provide teacher with understanding of what is going on.

Tutoring Techniques

- Heuristic Rewarding
- Decomposition
- Simplification
- Situation Selection
- Teleoperation
- Demonstration
- Coaching
- Explanation
- Cooperation
- Socratic method

Heuristic Rewards

- Giving the learner intermediate feedback about performance
- Related:
  - Reward shaping
  - Gamification
  - Heuristics in e.g. minimax game playing
- RL example:
  - Different reward for positive/negative step

Decomposition

- Decomposition of whole, complex tasks into smaller components

- Related:
  - Whole-task vs. part-task training
  - Curriculum learning
  - (Catastrophic interference)
  - (Transfer learning)
  - (Multitask learning)

- RL example:
  - Sliding puzzle at goal location on grid

Simplification

- Starting with a simplified version of the final task and gradually increasing the complexity
- Related:
  - Shaping (B.F. Skinner)
  - Curriculum learning
  - Decomposition
- RL example:
  - Slowly increase grid size and/or start near goal

Situation Selection

- Selecting situations (or data) for the learner to focus on
  - e.g. simpler or more difficult situations
- Related
  - Boosting
  - ML application development
  - Big Data
  - Active learning / teaching
- RL example:
  - Start (or stop) in problematic states

Teleoperation

- Temporarily taking control of the learner's actions so they can experience them
  - Right level of abstraction
- Applications:
  - Tennis / golf / chess
  - Robot ping pong
  - Artificial tutor
- RL example:
  - Force good or random moves

Demonstration

- Showing the learner how to accomplish a task
- Requirements:
  - Desire to imitate
  - Ability to map tutor's actions onto own actions
  - Generalization ability
- Related:
  - Apprenticeship learning
  - Inverse reinforcement learning
  - Imitation learning

Tutoring by Demonstration

- Show the learner what to do
- Add tutor observation dimensions to state
- Requirements:
  - Generalization
  - Desire to imitate
  - Ability to map tutor actions to learner actions

- Tabular Q-learning agent
- Simple grid navigation task

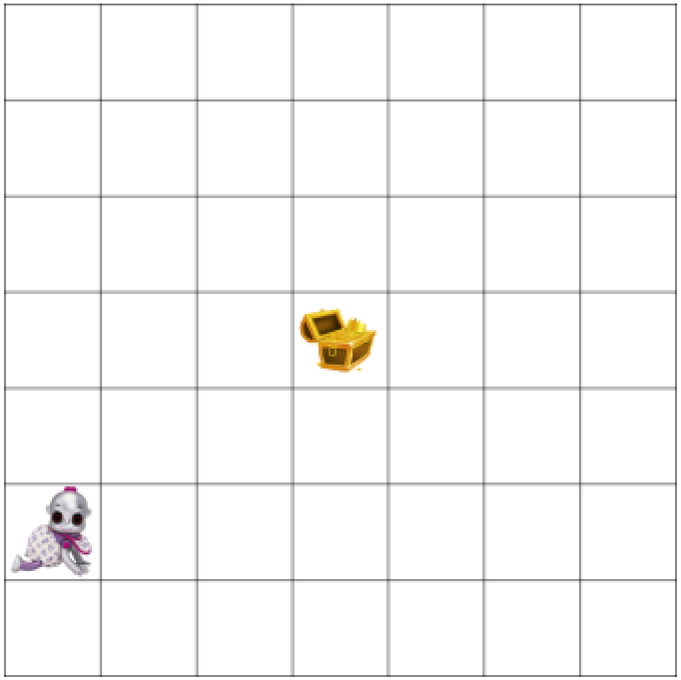

Coaching

- Giving the learner direct instructions of what action to take during the task
- Requirements:
  - Ability to map language-based instruction onto actions
  - Generalization ability
- Related:
  - Supervised learning
- RL example:
  - Add input that specifies correct output

Explanation

- Explaining to the learner how to approach certain situations before the starts (a new instance of) the task

- Requirements:
  - Language
  - Generalization ability

- Related:
  - Imperative programming
  - Analogies

Cooperation

- Doing a task together with the learner to facilitate other tutoring techniques

Socratic method

- Asking questions to encourage critical thinking and guide the learner towards its own conclusions.
- Related:
  - Shaping/chaining
- NARS example
  - <dog --> mammal>.
  - <<$x --> mammal> --> <$x --> [breaths]>>.
  - <{Spike} --> dog>.
  - <{Spike} --> [breaths]>? // main question
  - <{Spike} --> mammal>? // helping question

Tutoring Techniques

- Heuristic Rewarding
- Decomposition
- Simplification
- Situation Selection
- Teleoperation
- Demonstration
- Coaching
- Explanation
- Cooperation
- Socratic method

Artificial Pedagogy

- Systematization and theory of teaching artificial learning systems.
  - Focus on teaching rather than learning
  - Aimed at full spectrum of learning systems
  - Emphasis towards AGI-aspiring systems

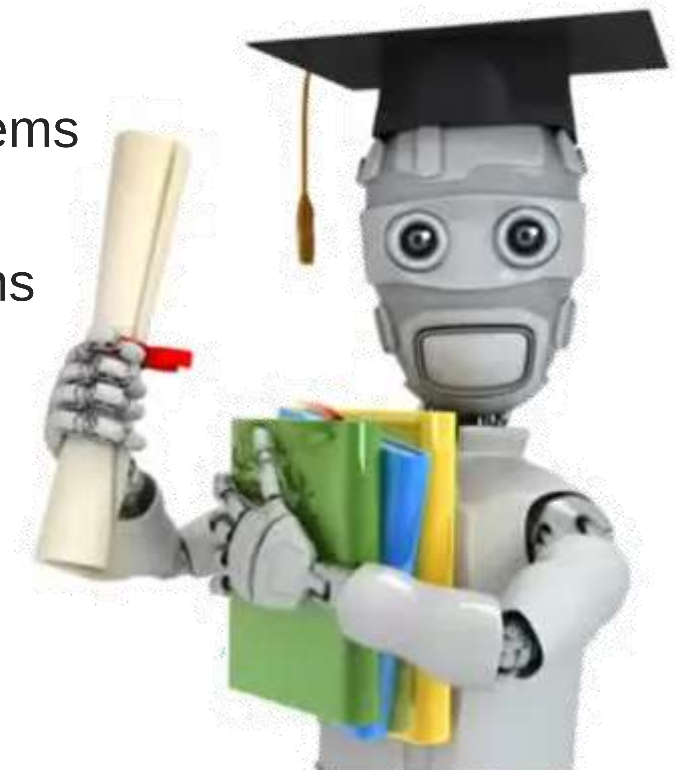

Why Artificial Pedagogy?

- No general theory exists
- Sophisticated teaching needed in complex domains
- Teaching is powerful but complex
- Current machine teaching is ad hoc
- Sufficiently advanced learners now exist
- Relevance will increase as field advances

Questions?