



Raising AI: Tutoring Matters

Jordi Bieger¹ (jbieger@gmail.com), Kristinn R. Thórisson^{1,2} & Deon Garrett²

¹Reykjavik University | Center for Analysis & Design of Intelligent Agents

²Icelandic Institute for Intelligent Machines

Path to Adult-Level AI



- Typical AI project:
 - The system only learns on the final task
 - The system is alone
- Raising AI:
 - Helping an AI system learn, grow from baby-AI into adult-AI, and realize its potential



Why raising?

- Guidance necessary to deal with complex new situations
- Less sophisticated system needed to reach the same level of intelligence
- Biologically plausible



Goals for the paper

- Argue for the importance of research into raising AI
- Discuss issues related to raising and tutoring
- Unite research from different fields under the perspective of raising AI
- Provide a starting point for various techniques for tutoring AI



Tutoring matters

- Focus on tasks rather than environments or cognitive stages
- Tutoring methods and learning algorithms impose requirements on each other
- Tutoring doesn't always help
- Tutoring can be difficult
- Human tutors may be expensive and/or inefficient



Tutoring Techniques

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



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



Questions?

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



end of presentation



Heuristic Rewards

- Giving the learner intermediate feedback about performance
- Related:
 - Reward shaping
 - Gamification
 - Heuristics in e.g. minimax game playing



Decomposition

- Decomposition of whole, complex tasks into smaller components
- Related:
 - Whole-task vs. part-task training
 - Curriculum learning
 - (Catastrophic interference)
 - (Transfer learning)
 - (Multitask learning)



Simplification

- Starting with a simplified version of the final task and gradually increasing the complexity
- Related:
 - Shaping (B.F. Skinner)
 - Curriculum learning
 - Decomposition



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



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



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



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



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

