

Developing a self-learning AI for a turn-based strategy game

Master's Thesis project proposal

Mikael Kalms, Fall Damage Studio AB
kalms@falldamagstudio.com

Brief

Fall Damage, a video games company based in Stockholm, are developing a turn-based strategy game. The game will offer one-versus-one play between humans. Fall Damage are interested in exploring the feasibility of constructing an AI for their game which can learn to play the game well with minimal human instruction. If successful, this may result in building the foundation of automated tests that are used during continued development of the product.

Background: AI research, machine learning and turn-based strategy games

AI research defines itself as the study of "intelligent agents". It is today an umbrella term that spans many different subfields - vision, natural language, decision theory, genetic algorithms, robotics etc.

Machine learning is an active AI research area which focuses on techniques for finding solutions to particular problems given various levels of instruction. Some approaches are based on finding structure from a large body of examples with different classifications; others are based on trying different choices out and then then being told how good the choices were. Video game AI is well suited to machine learning since the game simulation is always available in a digital representation.

Turn-based strategy games originate from games such as chess. There is a world, with a well-understood state; a set of clearly defined rules; at certain points the players are allowed to choose between a predefined set of actions and thereby affect the continued evolution of the world state; there is a well-defined winning condition. The well-defined rule set, time-discrete nature and comparably low number of turns per game makes it feasible to apply many techniques from AI research to the problem of constructing effective AI players. Good results require careful choice both of knowledge representations as well as planning strategies.

Project proposal

Fall Damage invites a student to investigate effective methods for creating effective AI opponents for their upcoming game. Most of the available literature analyzes well-established games such as Chess, Othello and Go. Game design differences make it non-obvious which of the existing methods still apply to this game though.

The focus of the project is to use machine learning techniques to learn how to play the game effectively. The game simulation can be run at an accelerated pace, and there is plenty of computational power available. However, there is little to no human player data available, and results are much more useful if it is available in hours or days (rather than months). Since the game will change during development, the AI is more useful the more often it can be re-trained.

The game is still in development. It is being built inside-out, so there will at all times be a playable core of the game. The boundaries and the set of available actions will change over time. The thesis project can choose to either follow the development of the game, or to take a snapshot of the game at one point in time and develop the techniques based on what is available in that snapshot.

Aim of Master Thesis

Explore different machine learning based methods for creating an effective AI opponent for the game. Draw on previously published results as well as literature in the field and distill a set of strategies that work for the particular turn-based scenario that the game offers. Choose a set of strategies that can be adapted as the game develops. If applicable, showcase results and stimulate further dialogue about testing turn-based games within the games community.

How to apply

Send an email to thesis@falldamagestudio.com with a CV and personal letter, and subject line "Developing a self-learning AI for a turn-based strategy game". You may be asked to take part in a technical interview and/or do a practical work test.

Practical details

Fall Damage are located in Stockholm. The thesis project will be pursued remotely, with regular video conferences and short on-site visits. Depending on the university, it may be possible for two Master's students to collaborate on this thesis project.

Related publications

Introduction to Machine Learning, Third Edition, by Ethem Alpaydin, ISBN 9780262325738

Reinforcement Learning: An Introduction, By Richard S. Sutton and Andrew G. Barto, ISBN 9780262332767

Playing Atari with Deep Reinforcement Learning, by Volodymyr Mnih et al, *NIPS Deep Learning Workshop*, 2013

Mastering the game of Go with deep neural networks and tree search, by David Silver et al, *Nature* 529, 2016