

EXAMENSARBETE Optimizing Reinforcement Learning Algorithms Using Design Of Experiments

STUDENT Anton Fristedt

HANDLEDARE Linda Hartman (LTH)

EXAMINATOR Elin A. Topp (LTH)

Robotics on Steroids: Optimizing Performance with Design of Experiments

POPULÄRVETENSKAPLIG SAMMANFATTNING **Anton Fristedt**

Do you ever find yourself in a situation where you try to improve something through trial and error? Maybe you want to make better bread, refine workout routines, grow healthier plants or perhaps optimize reinforcement learning algorithms used in robotic learning. But after a few attempts of improving, you quickly realize that you have no idea what you are doing. I have heard enough, you need more design of experiments in your life.

If the preamble was not obvious enough, this work explored the effectiveness of utilizing design of experiments to optimize reinforcement learning algorithms.

Reinforcement learning is a type of machine learning where a computer program learns how to solve a problem by repeatedly trying different solutions and receiving feedback on their success or failure.

In recent years, reinforcement learning has emerged as a highly effective approach in applications such as robotics and game playing. However, there is no one-size-fits-all solution in reinforcement learning, and finding the optimal setting for a given problem remains a key challenge. As such, the ability to fine-tune algorithms to match the specific task at hand should not be underestimated.

By applying design of experiments we observe how different changes to inputs causes changes in the output. When conducting a survey, ex-

periments are planned in advance and then analyzed with statistical methods. The result is easy to interpret and provides strong statistical evidence. This allows us to study multiple different input variables simultaneously and determine which ones that has a significant effect on the measured output value.

Our results showed that using design of experiments was successful in improving the performance of the reinforcement learning algorithm. Significant input variables was found and optimized, then evaluated against a model with default settings. The outcome of this evaluation demonstrated that our model reached the desired threshold of success 2.85 times faster. Is it perfect? Maybe not, improvements can always be made but it goes to show that with relative few experiments and simple methods, higher performance can be achieved. Good enough goes a long way, experimental design is applicable more or less everywhere and could benefit research in many fields.