Simulating the visual system of a deep reinforcement learning agent in different ways on Atari games

When training an agent on Atari games not all information on the screen is relevant. In this thesis the agent is tasked with controlling an area of interest that effectively can be used to exclude irrelevant pixels. Different types of interest areas are proposed, some of which perform well with limited information.

In reinforcement learning an agent is tasked with taking actions in such a way that some rewards are maximized. A reinforcement learning agent learns through trial and error; if the agent performs well it receives a positive reward and if the agent performs badly it receives a negative reward. In deep reinforcement learning these ideas have been combined with deep learning and neural networks. This combination allows reinforcement learning to be used on larger environments and on more complex problems, making reinforcement learning more useful and a more popular choice for solving problems.

When it comes to evaluating the performance of deep reinforcement learning agents different Atari 2600 games are often used with the goal of learning to play the game. These games are useful as the problem solving ability of an agent can be tested quite easily on them. The Atari games are, however, fully observable: all the information on the screen is available to the agent at all times. A problem with this is that not all this information is relevant to the agent. In an attempt to exclude irrelevant pixels, the agent was in this thesis given a small region of interest (ROI) which could be moved across the screen. With this the agent had to learn where too look at the same time as it had to learn how to play the game.

Throughout the thesis the agent's "visual system" for controlling the region of interest was simulated in different ways. First, the ROI was created using a simple rectangle containing pixels of one chosen resolution. The pixels outside this rectangle were simply replaced with zeroes. By using ROIs of different sizes and different resolutions it could be shown that a larger ROI with lower resolution was preferable to a smaller ROI with a higher resolution. The agent's field of view is, in other words, more important than the resolution.

Next, a new ROI was created by letting the resolution decrease as the distance to the center of the ROI increased. This type of ROI is inspired by the human visual system. Keeping the size of the ROI fixed while the resolution in the center was increased and the resolution towards the edges was decreased, did not, in general, improve performance. The increased resolution towards the center was probably not enough to make up for the lowered resolution towards the edges of the ROI. It was, however, clear that the low resolution towards the edges did provide important information to the agent.

A final addition to the models used in the thesis was peripheral vision. This addition is also inspired by the human visual system. The agent's peripheral vision was simulated using a very low resolution background which added very little information to the model. When using peripheral vision, information from the whole screen is used, even though the information given by the peripheral is too blurry to make out details. The very low resolution of the peripheral is, however, enough to track movements and detect changes on the screen. This is made clear by the fact that the addition of the peripheral, in general, significantly improved the performance in the games.

Throughout the thesis it has been shown that an agent can perform well while controlling its ROI and also that it is possible to reduce the amount of information available to the agent quite a lot without hurting the performance. The successful agents from this thesis could possibly be modified to work on more complex real-world environments, where the use of a visual system is needed. The different ROIs using a reduced amount of information could also be useful in situations where there are limits to bandwidth or something similar.