


# Playing Atari Games with Deep Q-learning and Attention

In the past years, a popular way to evaluate the performance of Reinforcement Learning-based agents has been through a suite of Atari 2600 games. This suite has famously been used to benchmark the progress of building successively ever more intelligent agents.

The famous reinforcement learning (RL) experiments on the Atari games use the complete image of the video game for learning which makes the state space unnecessarily large.

In this thesis we want to explore if we can improve RL by using some type of *focus-of-attention* mechanism that would allow the agent to dynamically focus on the part of the screen (state-space) where it is most important.

We will start out with a reimplementation of Deep-Q Learning to play the Atari games. In a second step we will aim to reduce the state space through a simple focusing mechanism. This approach is called *active vision*. In the Fig. below you can see the game Space Invaders, and a white box marks example focus areas.



Here, we can make use of results from an existing master thesis. Finally, we will push the boundaries by exploring different types of sensor resolutions that can further reduce the state space, e.g. a *log-polar* resolution that is motivated by the human eye. It has the highest resolution in the middle with a linearly reduced resolution with increasing distance from the center.

**To do this work, you should ideally have good python skills and ideally you have completed the applied machine learning course with a good grade. A course in computer vision can be good but is not a must.**

## Contact:

Volker Krüger	Professor i Datavetenskap	<a href="mailto:volker.krueger@cs.lth.se">volker.krueger@cs.lth.se</a>
---------------	------------------------------	--