



Remember these?



A large, white, stylized Atari logo is centered on the left side of the slide. It consists of three vertical bars of varying heights and two curved lines that sweep upwards from the bottom, forming a shape reminiscent of the Atari logo.

Playing Atari Games using RL

VARSHA LALWANI

AKSHAY MASARE

Motivation

May be we can design game players for each one of them!

But, how about an AI agent who can **learn** to play them **all**!

This is where the concept of a *general game player* come into the picture.

In this project we are trying to implement a deep reinforced learning based agent to play multiple video games.



Problem Statement

Learning to play *Breakout* using a convolutional neural network model trained with a variant of Q-learning, whose input would be raw pixels and whose output would be a value function estimating future rewards.



Concepts Involved

Reinforcement Learning

Q-Learning

Convolutional Neural Network



Reinforcement Learning and Q-Learning

In a reinforcement learning model, an agent takes actions in an environment with the goal of maximising a cumulative reward.

Q-learning is a model free form of RL

Algorithm:

Initialize $Q(s, a)$ arbitrarily

Repeat (for each episode):

Initialize S

Repeat (for each step of episode):

Choose a from s using policy derived from Q

(e.g. ϵ -greedy)

Take action a , observe r, s'

$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$

$s \leftarrow s'$

until s is terminal



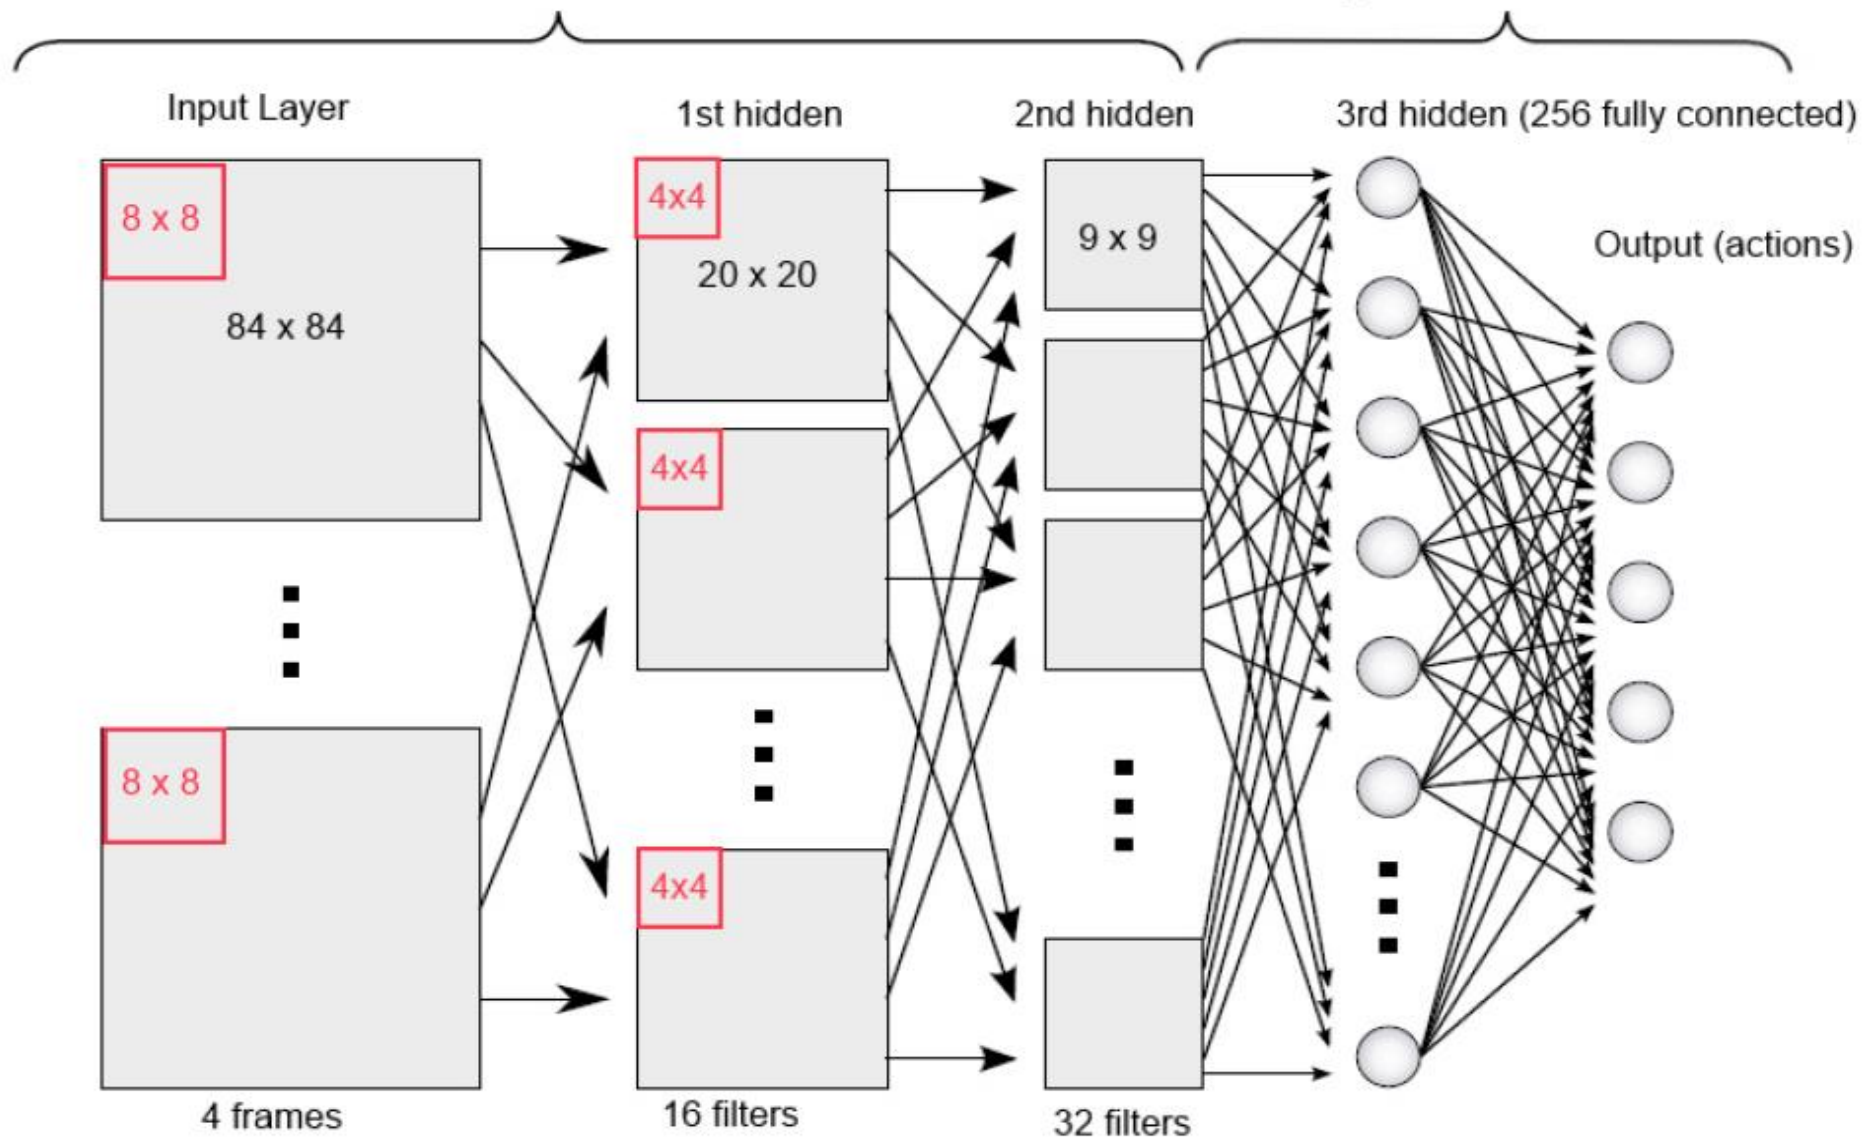
Convolutional Neural Networks

- Suited for extracting features from images
- We take 4 images at a time, downscaled to 84x84 pixels
- Images taken as 2D matrices
- 2D matrices convolved with linear filters
- Weight matrices for multiple image



Convolution

Fully connected



Nodes: 84 x 84 x 4

20 x 20 x 16

9 x 9 x 32

256

4

Weights:

8 x 8 x 4 x 16

4 x 4 x 16 x 32

9 x 9 x 32 x 256

256 x 4

Arcade Learning Environment

- It is built on top of Stella, open-source Atari 2600 emulator
- Built in C++, Support for over 50 games
- Can programmatically input player commands
- Outputs Image of the game screen, score and the state of the game



References

[1] The Arcade Learning Environment: An Evaluation Platform for General Agents by Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling Journal of Artificial Intelligence Research 47, pp. 253-279, 2013.

[2] Stella Emulator: <http://stella.sourceforge.net/>

[3] Playing Atari with Deep Reinforcement Learning by Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller NIPS Deep Learning Workshop, 2013.





Any Questions ??