End to end supervised driving in a basic OpenAI simulator

Motivation and objective

- Simulation is becoming a very useful tool for modeling the world.
- Autonomy is among the most impressive demonstration of progress and at the same time, one of the most sexiest and challenging projects in AI and ML.
- The objective is to have a feeling of how far we can go without RL (vs RL part in project CS221). With the coinstrains of using end to end models, random plays and no RL.
- We will try to explore the different pros and cons of different types of end to end models.

Can we train a linear model for it?

First we built a simple predictor made with the argmax of 4 linear models, one for each action. And use 3 succesive frames. The models are trained on random plays. We use this first predicor that achieves around 160 000 on average as a baseline. This is to compare the average 33 000 points of the data it has been trained on.

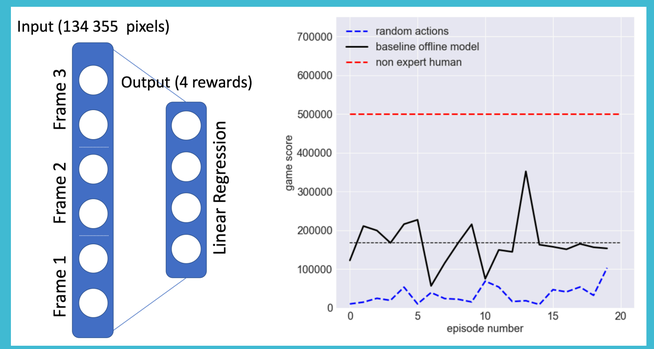

How about a fully connected neural net?

The baseline has the disadvantage of not learning the 4 actions at the same time. To addres this, and to try a more flexible model, we use a fully connected neural net. We mask the output nodes that don't correspond to the action taken in the training data and use MSE.

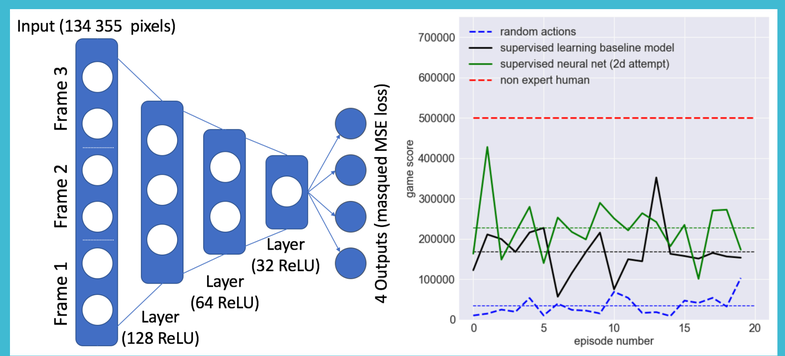

What if we used a convolutional neural net?

Quite disapointingly, this model should perform way better. Architectures of CNN turn out to be hard and as a beginer, it is just above the baseline

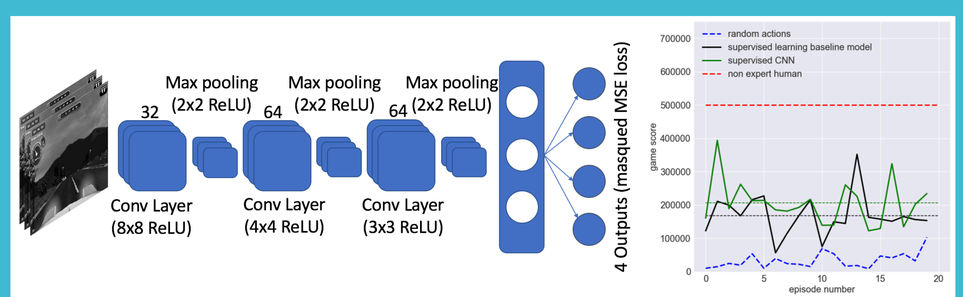

How well does an RNN?

We can see in the examples that speed can make much differences in the input images. Here the RNN seem to capture this well in the input sequence.

We used the same FCNN architecture but added a GRU cell in the First layer. This is our best performing model so far.

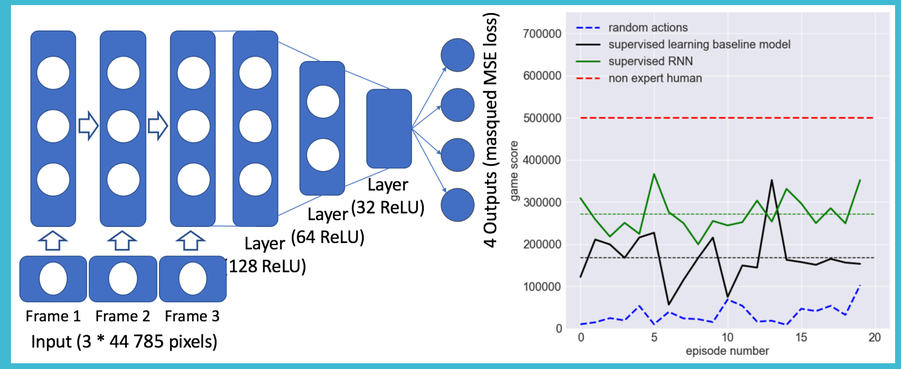

End to end model feature pipeline

Original image

Down sizing, grayscaling. 3 successive frames per state

Some results