




Playing Super Hexagon using Computer Vision

Seminar/Project Visual Computing
Final Presentation, March 9, 2021

Christoph Wiesmeier

Intro-Objective

- Let the PC play a Computer Game like a Human
 - Only Visual output
 - Input via Keyboard
 - Unmodified Game
 - In Real Time
- Super Hexagon




Link to Video [1]: <https://youtu.be/lzRgb-2Egp8>



Intro – Super Hexagon

- Scope of the Project
 - Three Implementation
 - Traditional CV
 - Supervised Learning
 - Reinforcement Learning
 - Compression of configurations
 - Prepossessing, Network Architechtues, ...
- 400h of Playtime
- Surpassing my playing skills


Processing (Initial)





Processing (Grayscale)




HSV Color-space → Value



Processing (Polar Transformation)



Polar Transformation



Processioning (Player Position)


Stack 2 Images

Find Player candidates


Filter/Select the player

x: [90,250]


y: nearest to center





Proceessioning (RadStab)



Traditional Computer Vision (Threshold)



Traditional Computer Vision (TCV)

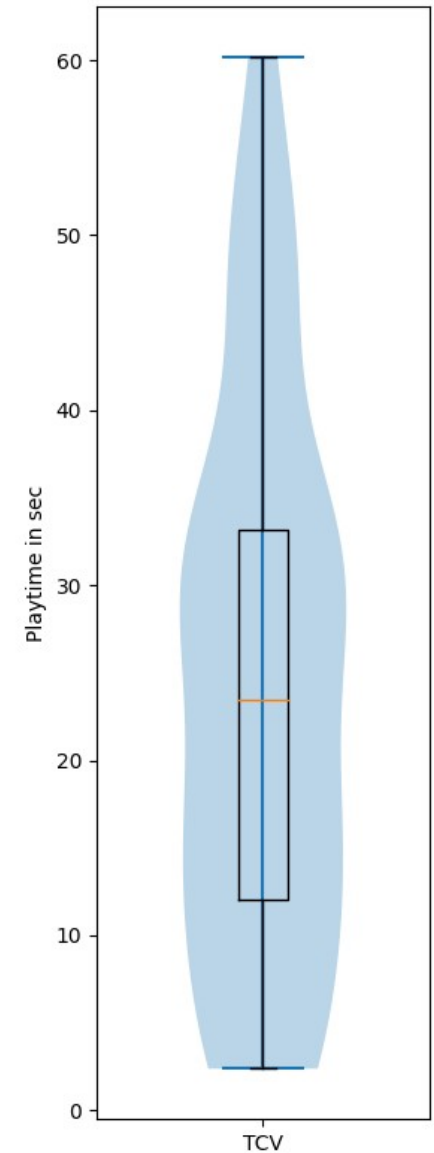


Traditional Implementation Results

- Results on 100 Rounds

| | Mean | Median |
|-----|----------|----------|
| TCV | 24.4 sec | 23.5 sec |

- Random Collisions
- Simplistic Model



Videos:

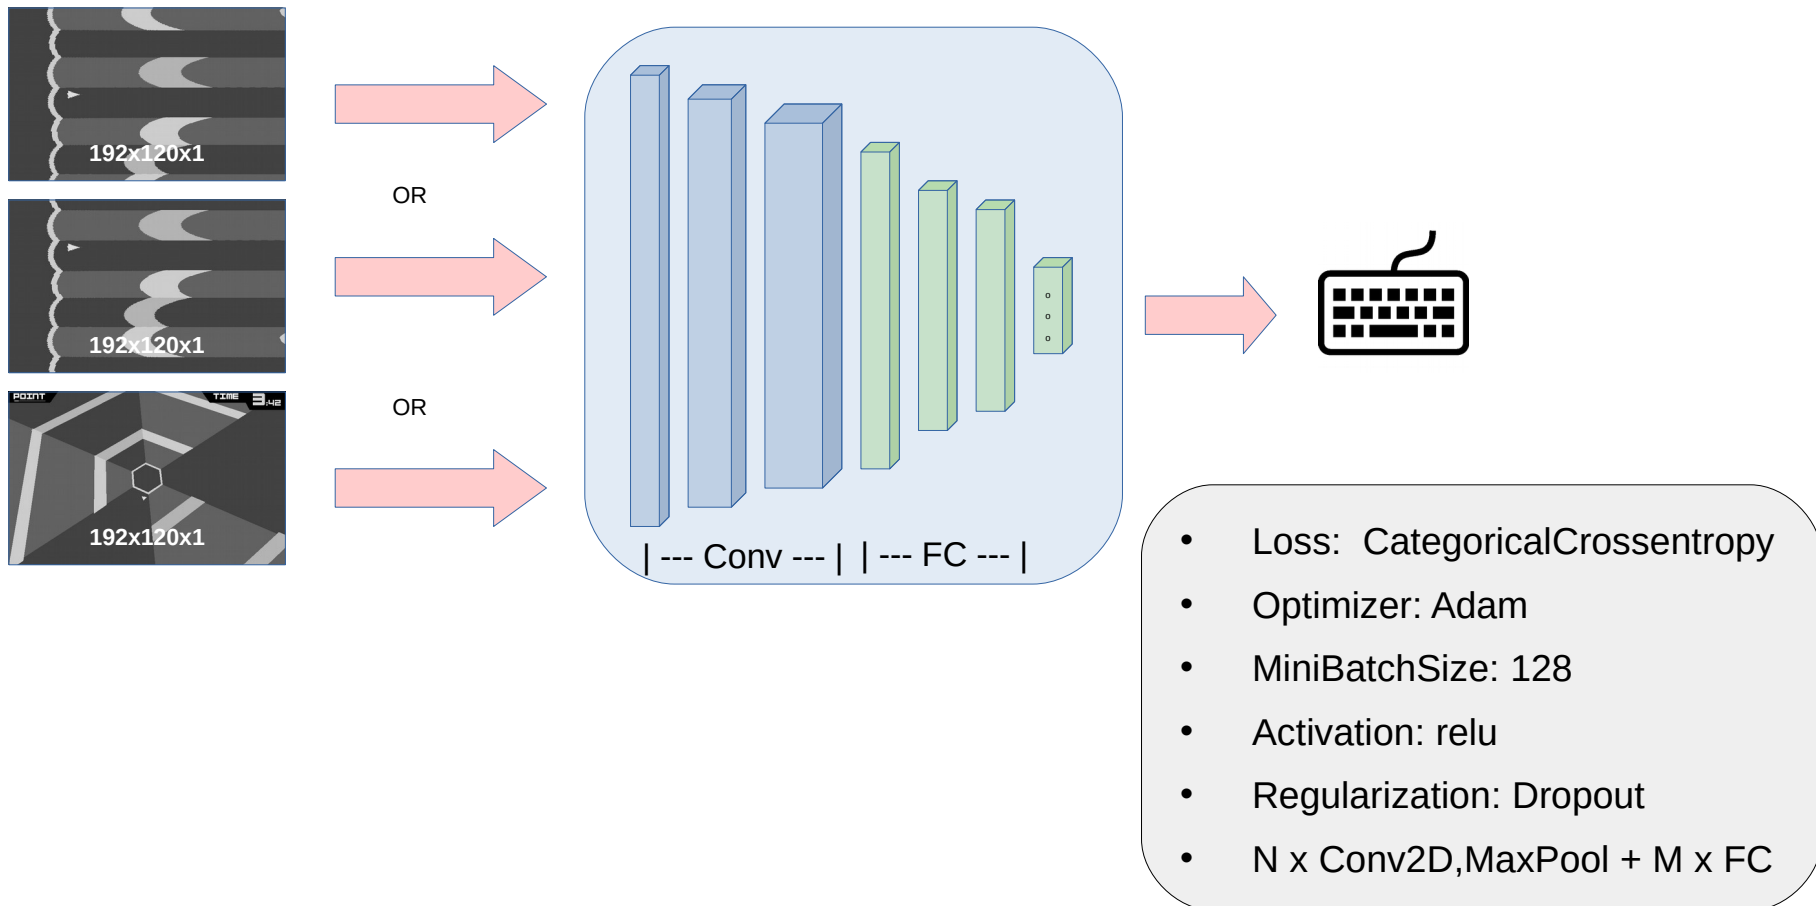
[2.1] <https://youtu.be/BJbC7AlodvU>

[2.2] <https://youtu.be/-8BWwtkEh0>

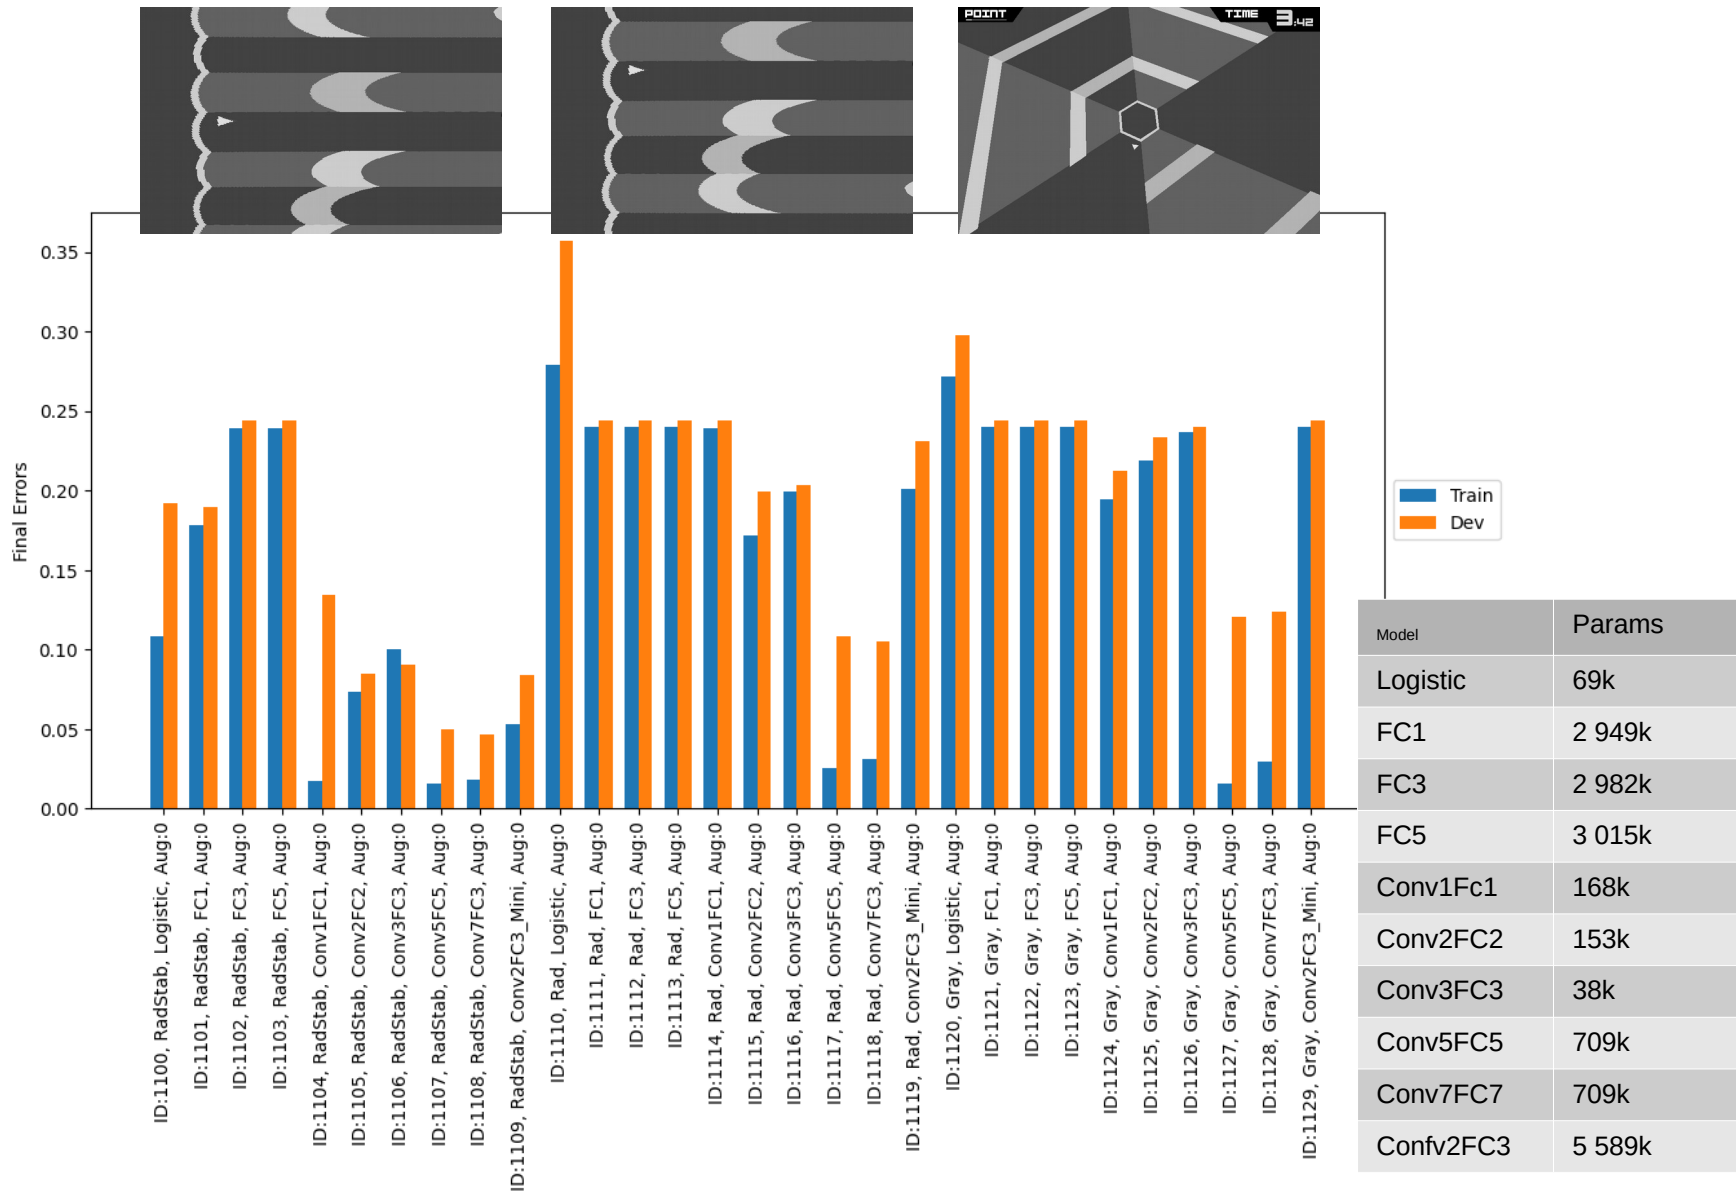
[2.3] https://youtu.be/oWF87jFd_XI (Traditional CV Internals)

Supervised Setup


- Implementation Based on Tensorflow/Keras
- Training Based on Human/TCV Data



Supervised Preprocessing/Network Size



Supervised Results




Video[3.1]: <https://youtu.be/p6qXnfXOrcM>

Video[3.2]: <https://youtu.be/bC6gQLMXSBY>

Reinforcement Implementation

- Algorithm: SARSA
- Reward: Only Crashes, Distance
- Processioning: Stabilized, Polar, Gray
- Function Approximation CNN (5Conv 4FC)
- Temporal Distance Learning:
 - 100ms
- Experience Replay

SARSA Introduction




- Value Method: $Q(\text{state}, \text{action})$

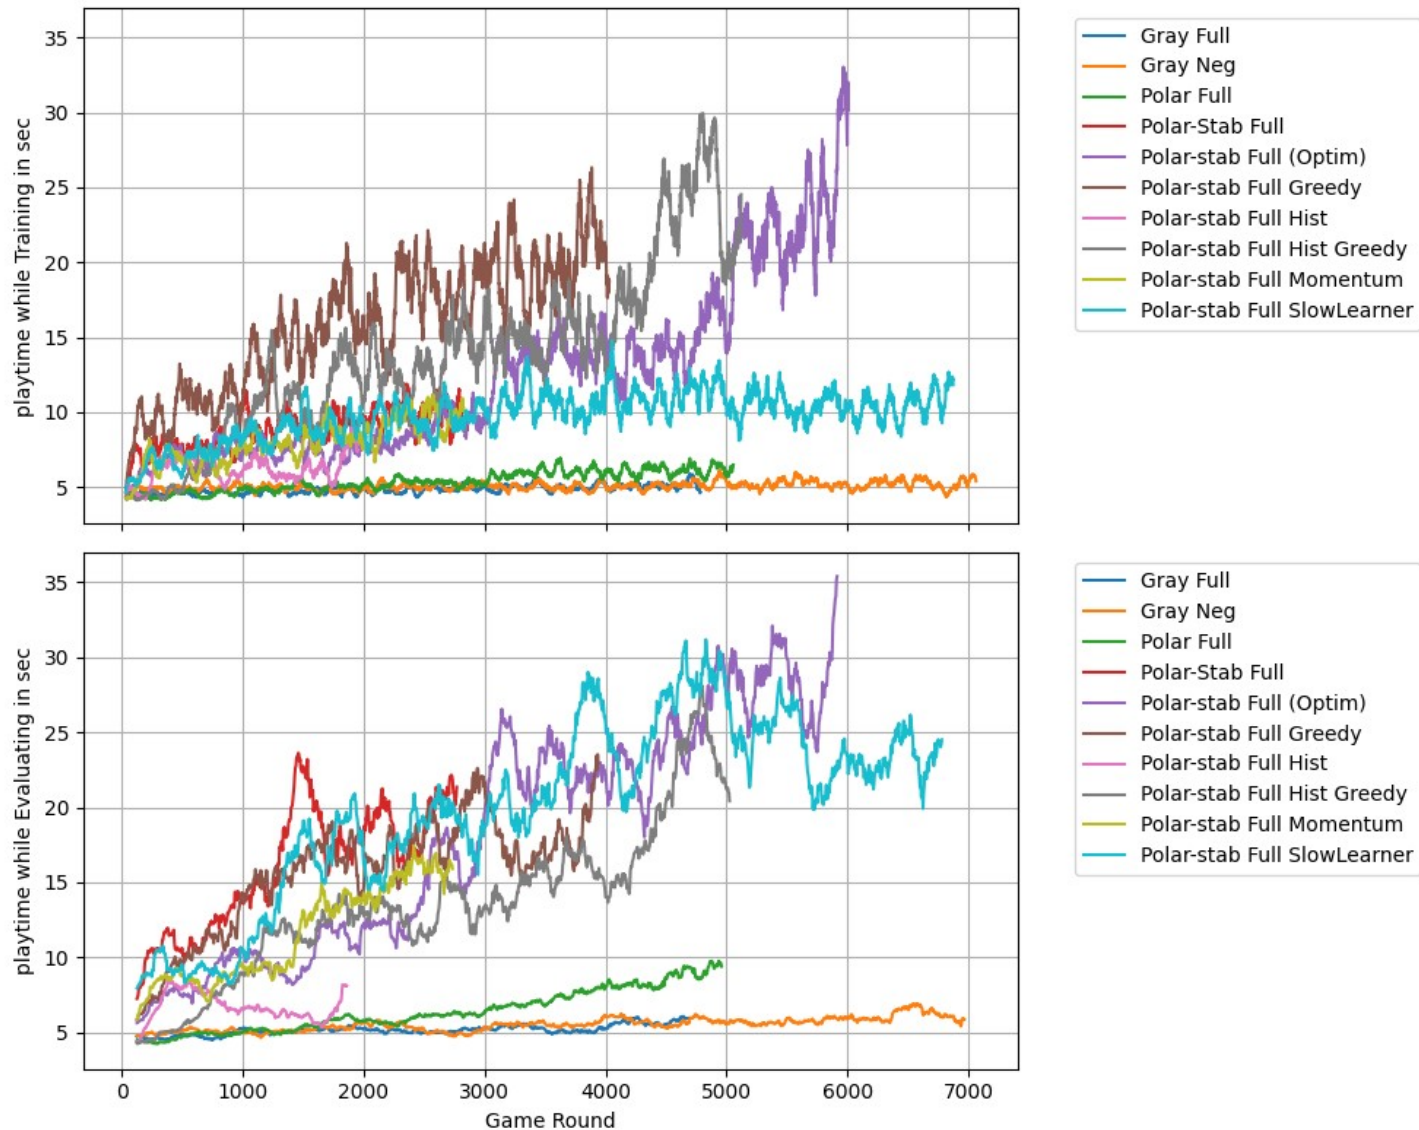
$$Q(s_i, a_i) \approx \sum_i \gamma^j * r_i$$

- Temporal Difference Learning

$$Q_{New}(s_i, a_i) = r + \gamma * Q(s_{i+1}, a_{i+1})$$



SARSA Results



Rolling mean over 50 samples

Comparison

- Human
- Traditional Computer Vision
- Supervised
- Reinforcement


| | Mean [sec] | Median [sec] |
|------------|------------|--------------|
| Human | 22.9 | 19.7 |
| TCV | 24.4 | 23.5 |
| Supervised | 30.7 | 30.6 |
| SARSA | 46.9 | 47.3 |

Video 4.1 <https://youtu.be/ySNdvXkl0qs>

Video 4.2 <https://youtu.be/v6LMM8U4ExQ>

Video 4.3 <https://youtu.be/4YUoSEf7J5w>

Video 4.4 <https://youtu.be/aRcwHBU0neE>



Conclusion

- Traditional Computer Vision
 - A lot of Tuning needed and cases to handle
- Supervised Learning can reproduce training results
- Reinforcement
 - Takes a lot of training
 - Image processing pipeline has big benefits



Playing Super Hexagon using Computer Vision

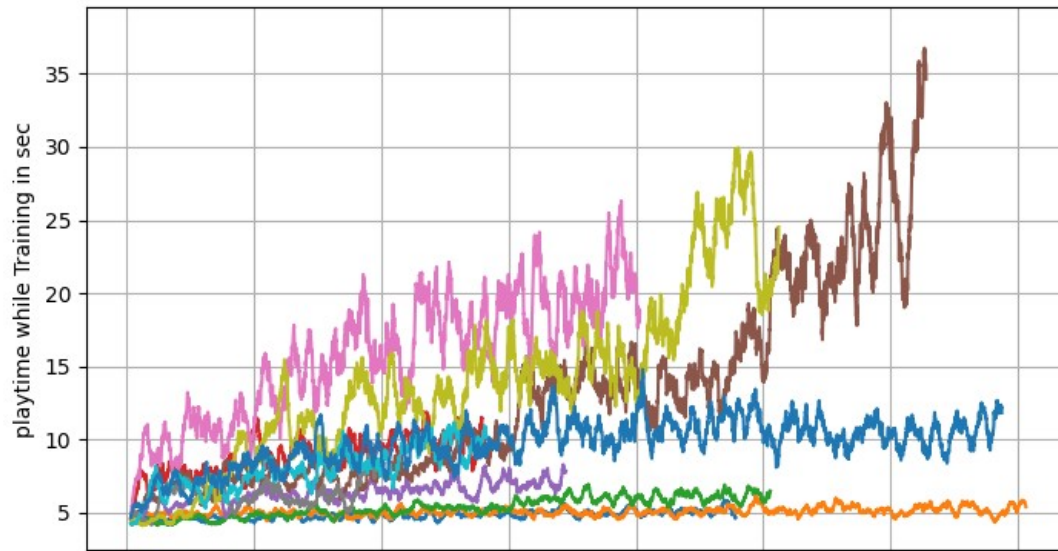
Seminar/Project Computer Vision WS 2020
Final Presentation, March 9, 2021

Christoph Wiesmeier

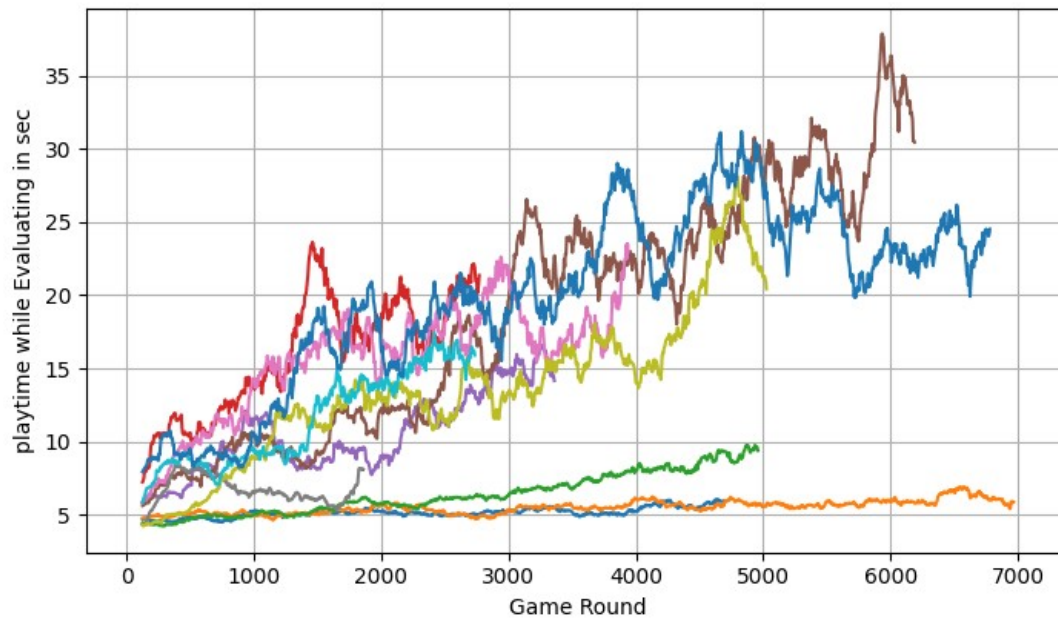


Videos

- [1] The Game (Human Play)
<https://youtu.be/lzRgb-2Egp8>
- [2.1] Traditional CV Implementation
<https://youtu.be/BJbC7AlodvU>
- [2.2] Traditional CV Implementation
<https://youtu.be/-8BWwitkEh0>
- [2.3] Traditional CV Implementation Debug
https://youtu.be/oWF87jFd_XI
- [3.1] Superviced
<https://youtu.be/p6qXnfXOrcM>
- [3.2] Superviced
<https://youtu.be/bC6gQLMXSBY>
- [4.1] SARSA configuration 12p
<https://youtu.be/ySNdvXkl0qs>
- [4.2] SARSA configuration 12p
<https://youtu.be/v6LMM8U4ExQ>
- [4.3] SARSA Configuration 13 (Average Run)
<https://youtu.be/4YUoSEf7J5w>
- [4.4] SARSA Configuration 13 (Best eval Run)
<https://youtu.be/aRcwHBU0neE>




- Gray Full
- Gray Neg
- Polar Full
- Polar-Stab Full
- Polar-Stab Neg
- Polar-stab Full (Optim)
- Polar-stab Full Greedy
- Polar-stab Full Hist
- Polar-stab Full Hist Greedy
- Polar-stab Full Momentum
- Polar-stab Full SlowLearner





- Gray Full
- Gray Neg
- Polar Full
- Polar-Stab Full
- Polar-Stab Neg
- Polar-stab Full (Optim)
- Polar-stab Full Greedy
- Polar-stab Full Hist
- Polar-stab Full Hist Greedy
- Polar-stab Full Momentum
- Polar-stab Full SlowLearner

SARSA Results





Traditional Implementation Limitations

The shorted way would be through the wall



Traditional Implementation Issues




Traditional Implementation Results

Issues:

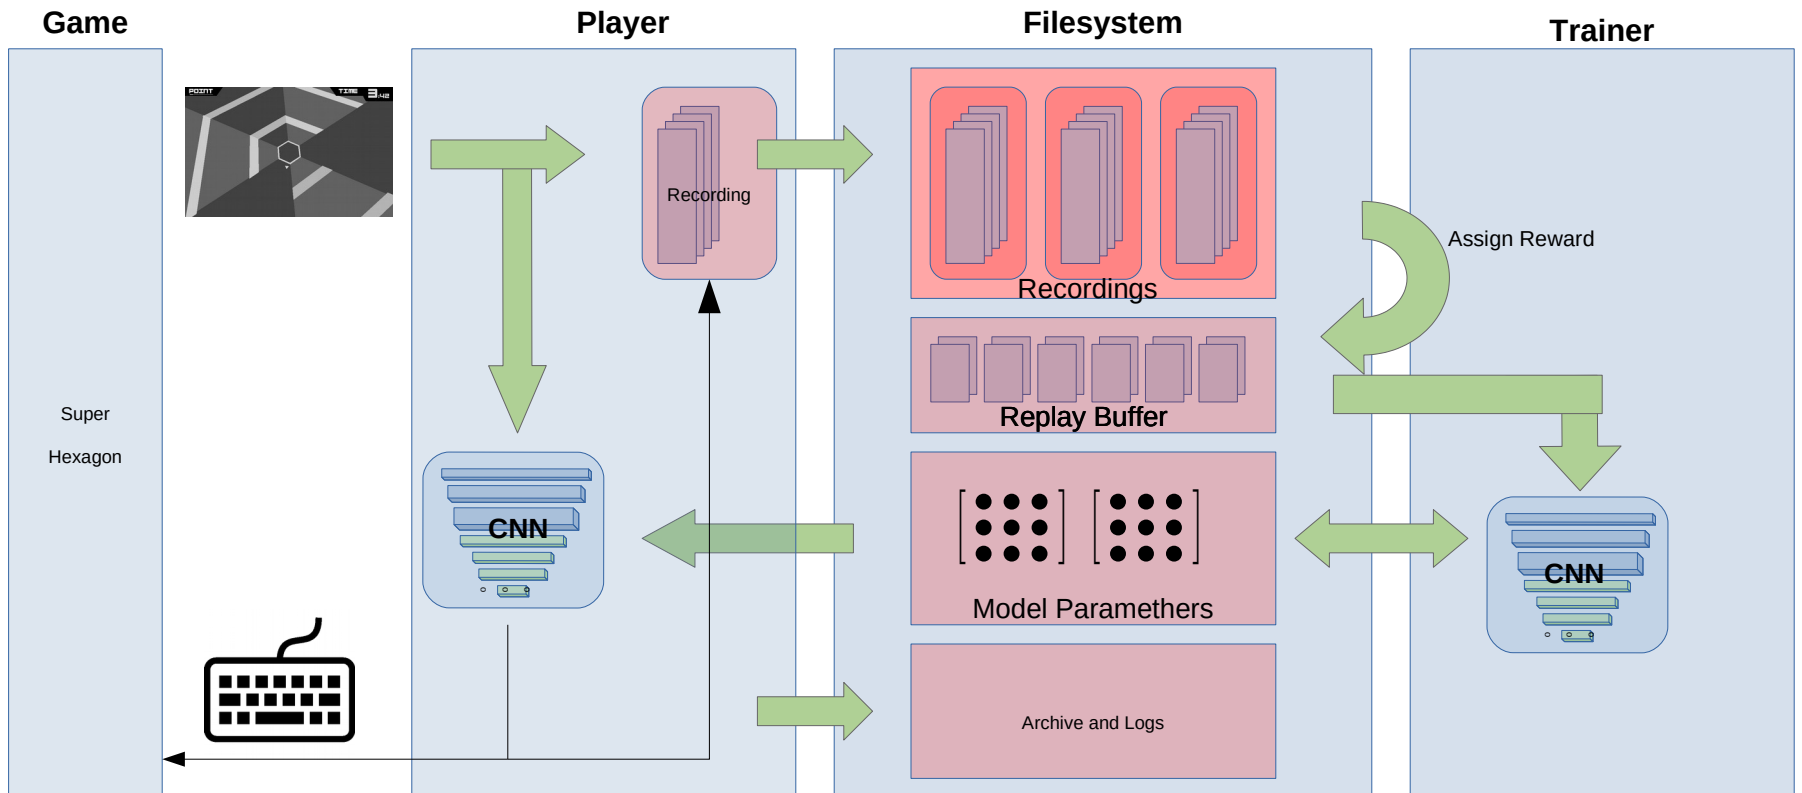
- Lag Related Errors (Approximately 3 Frames)
- No checking of free Path
- (Limited sight because of Polar Transformation)

Supervised Results



Reinforcement Architecture

- Algorithm: SARSA
- Reward: Only Crashes, Distance



Traditional Implementation

Processing Pipeline

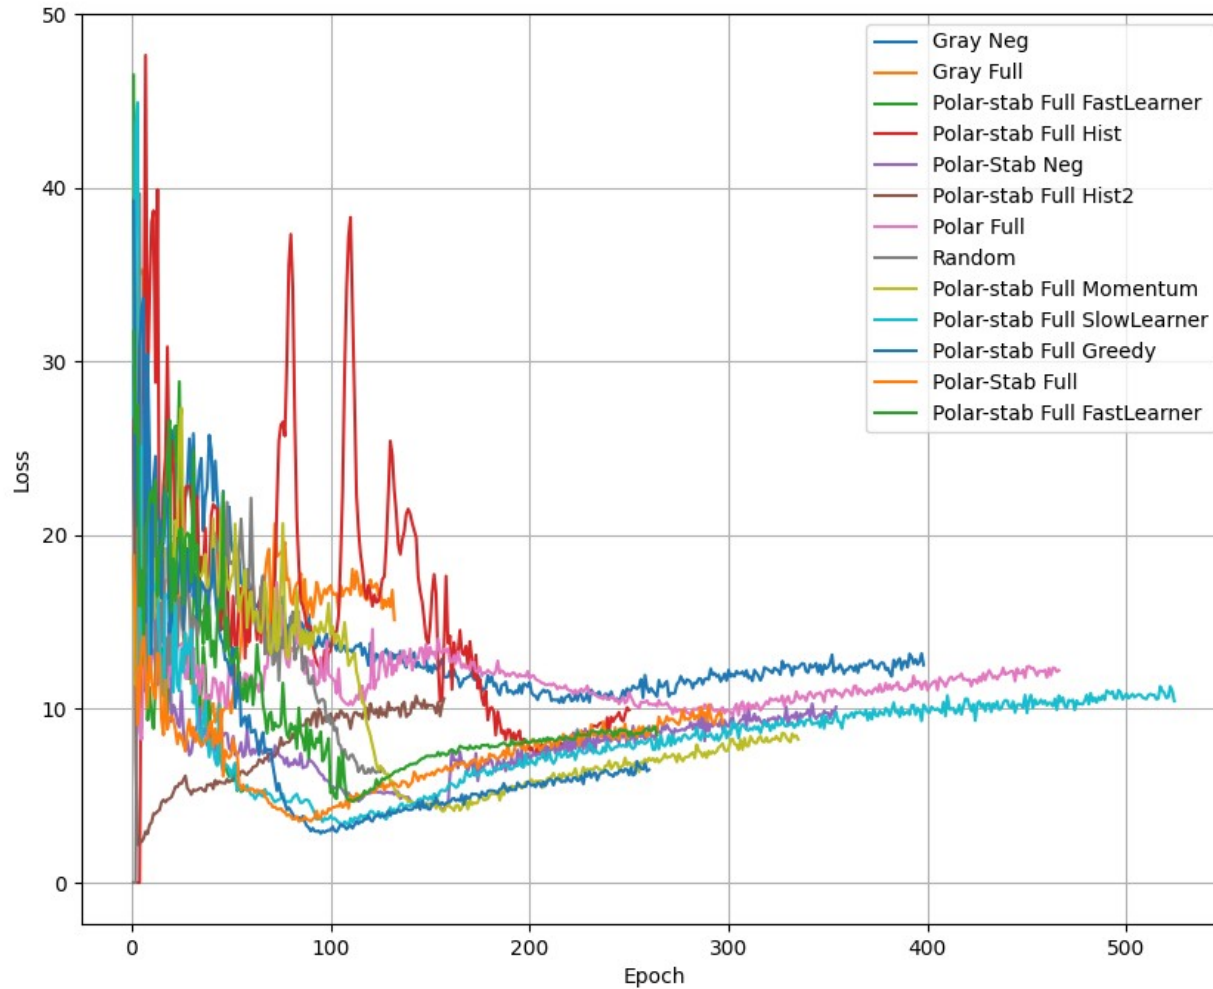
- Stabilize the Image
- Find the Target Position
- Execute move

- Suggestions on the Slide


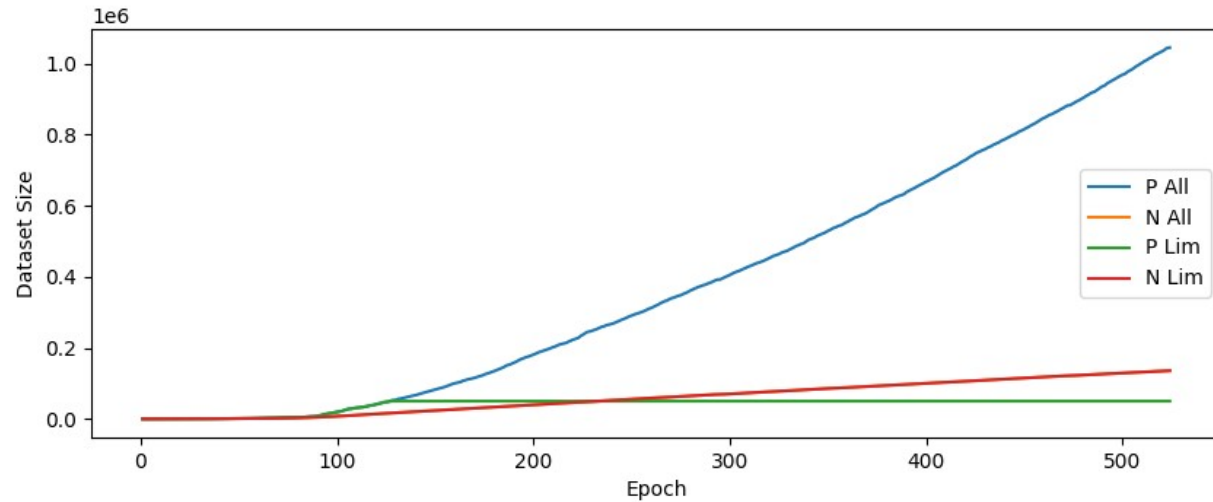
Notes

- Reinforcement Hyperparameters
 - TraingBufferSizes
 - Adam Momentum Reset/NoReset
 - Pos Reward Scaling
 - Training Buffer Weighting
 - Gamma
 - T-Delta
 - Models
 - FC,CNN,Depth,KernelSizes
 - Epsilon Scedule
 - Optimizer Params
 - Image Resoltion
 - Dueling Network

SARSA Loss



SARSA Dataset Size



Possible Additions

- Possible Additions
 - Show the Frame Debugger
 -

Links

- [1] <https://youtu.be/lzRgb-2Egp8> (Game Example)
- [2.1] <https://youtu.be/BJbC7AlodvU> (Traditional CV)
- [2.2] <https://youtu.be/-8BWwitkEh0> (Traditional CV)
- [2.3] https://youtu.be/oWF87jFd_XI (Traditional CV Internals)

Supervised data sets

Recorded Image → Used Command
0.1s time interval

| Name | Info | Samples Train | Samples Dev |
|----------|---------------------------------|---------------|-------------|
| Human A | About 1.5h Playtime | 22752 | 2527 |
| Human100 | 100 Rounds Human | 18481 | 2053 |
| TCV100 | 100 Rounds PC | 20418 | 2268 |
| Both | TCV100 + Human100 | 38898 | 4322 |
| Large | All above + ~500 TCV Recordings | 161232 | 17914 |

Template Matching

ESC - STAGE SELECT


HIGH SCORES

LAST





NEW RECORD

PRESS SPACE TO RETRY

TCV Parameter Tuning




Supervised Training Results



Reinforcement Implementation


- Each Image is Stored Once
- Reference to Prior State





Important to highlight

- Asynchronous/Realtime /No Lab style Environment
-



Supervised Model

```
Sequential([
  Conv2D(filters=3,kernel_size=5,strides=(1, 1),activation="relu",input_shape=inputShape),
  MaxPooling2D(pool_size=(2, 2),strides=(1, 1)),
  Conv2D(filters=3,kernel_size=5,strides=(1, 1),activation="relu"),
  MaxPooling2D(pool_size=(5, 5),strides=(3, 3)),
  Conv2D(filters=3,kernel_size=5,strides=(1, 1),activation="relu"),
  MaxPooling2D(pool_size=(5, 5),strides=(3, 3)),
  Flatten(),
  Dropout(0.2),
  Dense(64, activation='relu'),
  Dropout(0.2),
  Dense(32, activation='relu'),
  Dropout(0.2),
  Dense(16, activation='relu'),
  Dropout(0.05),
  Dense(3,activation="softmax")
])
```

- Loss: CategoricalCrossentropy
- Optimizer: Adam
- MiniBatchSize: 128
- Activation: relu
- Regularization: Dropout
- N x Conv2D,MaxPool + M x FC