A.L.V.I.N.N.

Autonomous Learning Vehicle Integrating Neural Networks

Project Description

Developed a system for detecting airborne aircraft in collaboration with Rockwell Collins using

- Computer Vision
- Machine Learning
- Neural Networks

Requirements

Functional Requirements

- Process a single image or a continuous video stream
- Detect multiple objects in one image frame
- Report confidence levels of identified objects

Non-Functional Requirements

- Performance
- Scalability
- Extensibility
- Accuracy
- Reliability
- Throughput

Design Approach

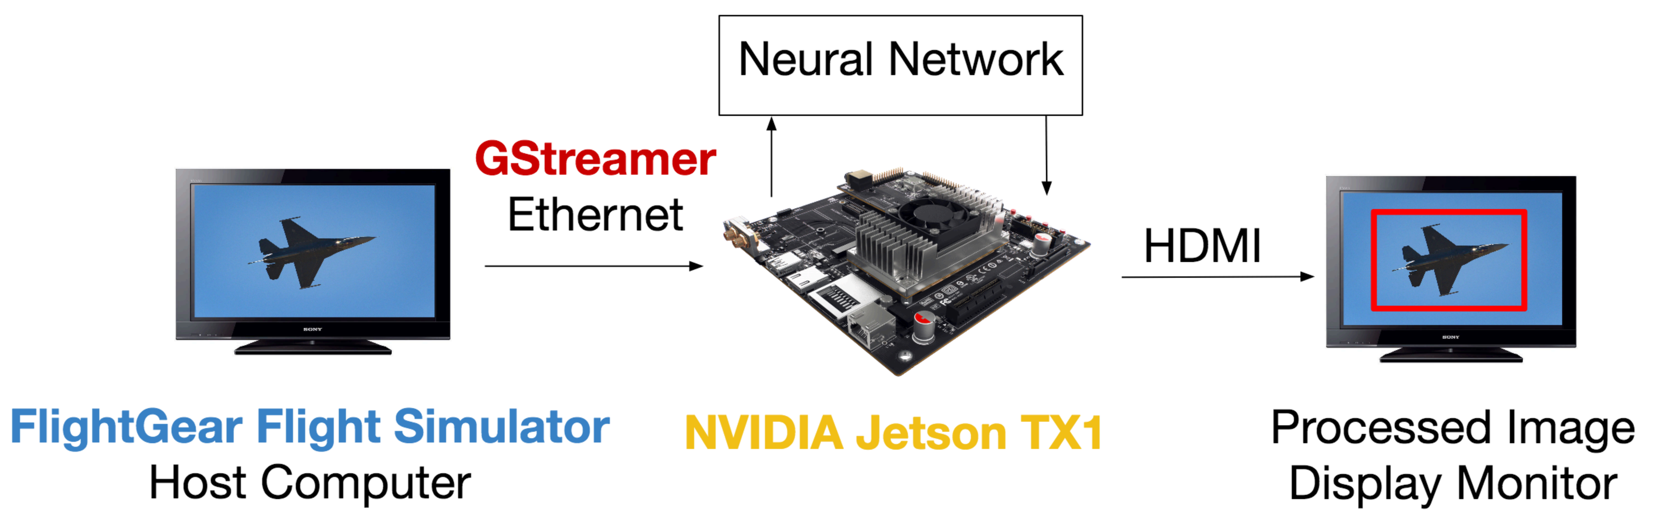

Flight Gear Flight Simulator

- Multi-platform open source flight simulator
- Used for video streaming and individual images

GStreamer

- Open source multimedia streaming application framework
- Streams the host computer's desktop to the embedded board

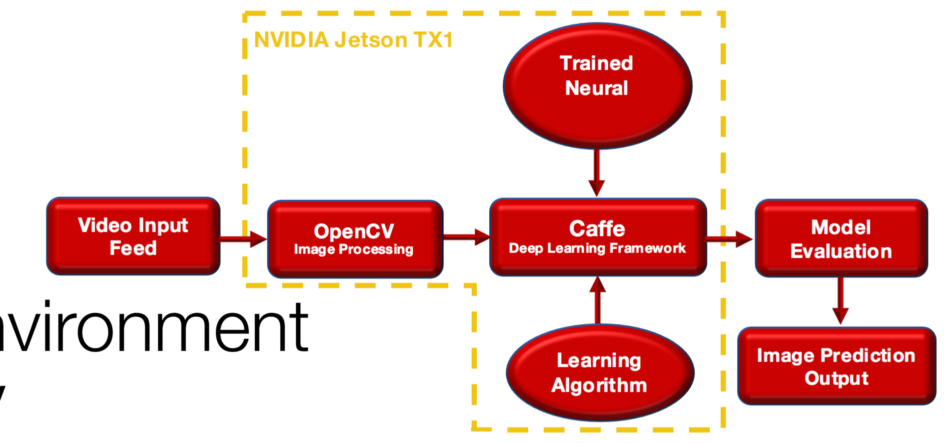

NVIDIA Jetson TX1

- GPU capabilities
- Linux operating environment
- Supports OpenCV
  - o Open source computer vision library
  - o Captures and resizes a frame from the feed
  - o Displays the output
- Supports Caffe
  - o Open source deep learning framework
  - o Trains and executes neural networks

Testing

Five different pre-made neural networks were tested and evaluated on their performance. Google MobileNets was tested using three different image sizes: 150x150 PX, 300x300 PX, and 450x450 PX.

Each network was tested on a set of 30 images containing plane and non-plane objects:

Performance Statistics (%)

Model	Sensitivity	Precision	Negative Predictor Value	Accuracy	Miss Rate	Fallout
DetectNet	76.92	88.24	52.63	75.47	23.08	28.57
SSD	56.41	100.00	45.16	67.92	43.59	0.00
YOLOv2	71.79	100.00	56.00	79.25	28.21	0.00
Tiny YOLO	46.15	100.00	40.00	60.38	53.85	0.00
MobileNets 150	25.64	100.00	32.56	45.28	74.36	0.00
MobileNets 300	66.67	100.00	51.85	75.47	33.33	0.00
MobileNets 450	71.79	100.00	56.00	79.25	28.21	0.00

The performance of each network under a continuous video stream was recorded:

Performance Metrics

Model	FPS	Memory Usage (GB)	Power (mW)
DetectNet	9.60	2.34	10120
SSD	5.57	2.57	11555
YOLOv2	1.75	3.23	11994
Tiny YOLO	6.20	2.21	10508
MobileNets 150	4.15	2.73	6678
MobileNets 300	3.56	2.84	7292
MobileNets 450	3.02	3.04	9166
Idle CPU	n/a	1.16	3343

Image Throughput per Watt

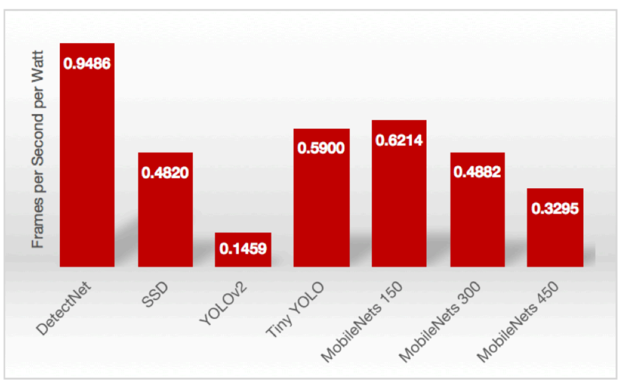

Results & Conclusion

Different models have different strengths and weaknesses, making them more applicable for certain operating environments:

- YOLOv2 and Google MobileNets 450 had greater accuracy towards detecting aircraft and the latter consumed less power; however, neither network performed well for real-time video processing
- DetectNet had the highest probability of detection, the lowest miss rate, and the best FPS performance, but it also had the greatest fallout
- Google MobileNets 300 represents a good choice regarding the power consumption, performance, and stability tradeoffs
- More data, tolerance modifications, and network retraining could yield improved performance

Meet the Team

Advisors: Drs. Jones and Zambreno
Client: Josh Bertram, Rockwell Collins

Bijan Choobineh (CprE)
Darren Davis (CprE)
Tracy La Van (CprE)
Jesse Luedtke (CprE)
David Schott (SE)
Robert Stemig (CprE)

ALVINN

For more information, visit our website: www.dec1709.sd.ece.iastate.edu