

Object Tracking Kick-Off

IEEE UAV Drone

J. Setpal

September 8, 2022



Table of Contents

- ① Low-Power Computer Vision
- ② Approaches
- ③ MLOps
- ④ Timeline
- ⑤ Final Thoughts

Table of Contents

① Low-Power Computer Vision

② Approaches

③ MLOps

④ Timeline

⑤ Final Thoughts

What does it mean?

What does it mean?

Resource Constrained.

What does it mean?

Resource Constrained. What resources do we have then?

What does it mean?

Resource Constrained. What resources do we have then?

- Intel x5-z8350
- Intel NCS2
- NU3000 depth Processor
- 4GB RAM
- 64GB ROM
- TOF (time-of-flight) Infrared Distance Sensor
- PMW3901 Optical Flow Sensor
- Depth Camera
- 160-degree Camera
- Inertial Measurement Unit

Not bad at all!

Low-Power Computer Vision

Low *Power* doesn't imply Low *Resources*.

Low-Power Computer Vision

Low *Power* doesn't imply Low *Resources*.

Our goal to maximize our resources.

Table of Contents

① Low-Power Computer Vision

② Approaches

③ MLOps

④ Timeline

⑤ Final Thoughts

Model Quantization

Model quantization involve floating-point weight conversions to lower-bit complements.

Model Quantization

Model quantization involve floating-point weight conversions to lower-bit complements.

Specialized hardware – like Intel NCS2 – are rapid at lower-bit computations.

Model Quantization

Model quantization involve floating-point weight conversions to lower-bit complements.

Specialized hardware – like Intel NCS2 – are rapid at lower-bit computations.

There's **post-model quantization**, which applies quantization after training. This is not as accurate as if training is conducted with quantization, called **quantization-aware training**.

Neural Architecture Search

Last year, we trained the **Nanodet** Architecture on our solution.

Neural Architecture Search

Last year, we trained the **Nanodet** Architecture on our solution.

It's a highly optimized solution, running at approx. 97fps on an ARM CPU. The general idea is to evaluate architectures on a case-by case basis to evaluate which fits best to the task.

Knowledge Distillation

Knowledge Distillation is the method of training a large, unwieldy network on a smaller neural network, to *distill* the knowledge contained within the larger network.




Table of Contents

- ① Low-Power Computer Vision
- ② Approaches
- ③ MLOps
- ④ Timeline
- ⑤ Final Thoughts

Maximizing Resource Utilization

Although Architectures like Nanodet are incredibly efficient, they aren't **task-specific**.

Maximizing Resource Utilization

Although Architectures like Nanodet are incredibly efficient, they aren't **task-specific**.

As a result, it fails to utilize information specific to our situation, which may be useful in developing an optimal solution for the competition.

This includes tools like the many sensors which we didn't use in our last year's solutions.

Simulation for Augmented Reinforcement-Based Training

Optimizing our dataset goes a long way towards developing functional models.

Simulation for Augmented Reinforcement-Based Training

Optimizing our dataset goes a long way towards developing functional models.

- Real data training works very well. However, it is nearly impossible to capture rare event.

Simulation for Augmented Reinforcement-Based Training

Optimizing our dataset goes a long way towards developing functional models.

- Real data training works very well. However, it is nearly impossible to capture rare event.
- Synthesized data doesn't work as well, but makes it possible to capture lots of edge cases.

Simulation for Augmented Reinforcement-Based Training

Optimizing our dataset goes a long way towards developing functional models.

- Real data training works very well. However, it is nearly impossible to capture rare event.
- Synthesized data doesn't work as well, but makes it possible to capture lots of edge cases.

The goal is to combine both!

Thank you!

Have an awesome rest of your day!

Slides:

<https://cs.purdue.edu/homes/jsetpal/slides/ot-kickoff.pdf>