



Deep Learning to Identify and Predict Objects in the Environment

Nicholas Sagers¹, Juan Andrade², Robert LiKamWa^{1,2}, Andreas Spanias².

¹ School of Arts, Media and Engineering (AME), Arizona State University

² School of ECEE, Arizona State University.



MOTIVATION

- Utilize advances in deep learning to assist in the ubiquitous tasks of memorization of object placement
- Employ advanced power savings techniques to drastically improve feasibility in low powered devices

PROJECT AIM

- This project will focus on the application of these models, the feasibility and implementation of deep learning tasks on low powered devices (Microsoft HoloLens), and future applications of this system.

PROBLEM STATEMENT


- To create a platform that assists individuals in locating objects in their immediate environments via computer vision and mobile systems

CHALLENGES


- Quickly detecting objects in the environment
- Doing so efficiently and reliably both in respect to client and server operations
- Leveraging available hardware to accomplish the task

METHOD


- YOLOv3: model for real time object classification
- Flask: client/server integration and secure data transfer
- NVIDIA CUDA: server side GPU Acceleration



REAL WORLD OBJECT DETECTION



SERVER RESULTS



QUALITATIVE RESULTS

- We found YOLOv3 to fit the project needs most effectively (although Tiny YOLO for serverless solutions may be more feasible, although the issue of accuracy still remains)
- The client software implementation used very little hardware resources, future implementations may include ways to leverage this hardware to expand existing capabilities

REFERENCES

- [1] M. Buckler, P. Bedoukian, S. Jayasuriya, and A. Sampson, "EVA2: Exploiting temporal redundancy in live computer vision," in 2018 ACM/IEEE 45th Annual Int. Sym. on Comp. Architecture, 2018, pp. 533–546.
- [2] Y. Zhu, A. Samajdar, M. Mattina, and P. Whatmough, "Euphrates: Algorithm-soc co-design for low-power mobile continuous vision," arXiv:1803.11232, 2018.
- [3] G. Evans, J. Miller, M. I. Pena, et al, "Evaluating the microsoft hololens through an augmented reality assembly application," in Degraded Environments: Sensing, Processing, and Display, vol. 10197.
- [4] X. Chen and A. Gupta, "An implementation of faster RCNN with study for region sampling," arXiv:1702, 2017.
- [5] F. N. Iandola, S. Han, M. W. Moskewicz, et al, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size," arXiv:1602.07360, 2016.

