

文章書き

だいたい切をのばして、待ちに待ってもらっている文章を書く仕事。

メモ

- OpenPose-Plus: <https://github.com/tensorlayer/openpose-plus>
- Detect-and-Track: Efficient Pose Estimation in Videos - https://www.zpascal.net/cvpr2018/Girdhar_Detect-and-Track_Efficient_Pose_CVPR_2018_paper.pdf
 - We propose an extremely lightweight yet highly effective approach that builds upon the latest advancements in human detection [17] and video understanding [5]
- Deep Learning Based 2D Human Pose Estimation: A Survey - <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8727761>
- Object Detection with Deep Learning: A Review - <https://arxiv.org/pdf/1807.05511.pdf>
 - In this paper, we provide a review on deep learning based object detection frameworks. Our review begins with a brief introduction on the history of deep learning and its representative tool, namely Convolutional Neural Network (CNN). Then we focus on typical generic object detection architectures along with some modifications and useful tricks to improve detection performance further.
- Object detection: speed and accuracy comparison (Faster R-CNN, R-FCN, SSD, FPN, RetinaNet and YOLOv3) - https://medium.com/@jonathan_hui/object-detection-speed-and-accuracy-comparison-faster-r-cnn-r-fcn-ssd-and-yolo
 - 後半は <https://arxiv.org/pdf/1611.10012.pdf> から
- Speed/accuracy trade-offs for modern convolutional object detectors - http://zpascal.net/cvpr2017/Huang_SpeedAccuracy_Trade-Offs_for_CVPR_2017_paper.pdf
 - Faster RCNN/R-FCN/SSD, Inception Resnet V2/Inception V2/Inception V3 /MobileNet/Resnet 101/VGG の精度, 処理時間, メモリ使用量の比較
 - <https://arxiv.org/pdf/1611.10012.pdf> - こっちの方がグラフがわかりやすい?
- 3D Human Pose Machines with Self-supervised Learning - <https://arxiv.org/pdf/1901.03798.pdf>
 - The main contributions of this work are three-fold. i) We present a novel model that learns to integrate rich spatial and temporal long-range dependencies as well as 3D geometric constraints, rather than relying on specific manually defined body smoothness or kinematic constraints; ii) Developing a simple yet effective self-supervised correction mechanism to incorporate 3D pose geometric structural information is innovative in literature, and may also inspire other 3D vision tasks; iii) The proposed self-supervised correction mechanism enables our model to significantly improve 3D human pose estimation via sufficient 2D human pose data.
- Ensemble Convolutional Neural Networks for Pose Estimation - https://www.toyota-ti.ac.jp/Lab/Denshi/iim/ukita/MyPapers/CVIU2018_Ensemble_preprint.pdf
 - In this paper, we present a PM-ensemble (PME) model to infer body configurations by modeling the interdependence among the responses of PM models.
- OpenPose - <https://modelzoo.co/model/openpose-caffe> <https://arxiv.org/pdf/1812.08008.pdf>
 - <https://github.com/CMU-Perceptual-Computing-Lab/openpose>
- YOLO TensorFlow - <https://modelzoo.co/model/yolo-tensorflow-tensorflow>
- Keras YOLOv3 - <https://modelzoo.co/model/keras-yolov3>
- Very Deep Convolutional Networks for Large-Scale Image Recognition - <https://modelzoo.co/model/very-deep-convolutional-networks-for-large-scale>
- Tensorflow detection model zoo - https://github.com/tensorflow/models/blob/master/research/object_detection/g3doc/detection_model_zoo.md
 - We provide a collection of detection models pre-trained on the COCO dataset, the Kitti dataset, the Open Images dataset, the AVA v2.1 dataset and the iNaturalist Species Detection Dataset.
- PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes - <https://arxiv.org/pdf/1711.00199.pdf>
 - In this work, we introduce PoseCNN, a new Convolutional Neural Network for 6D object pose estimation. PoseCNN estimates the 3D translation of an object by localizing its center in the image and predicting its distance from the camera.
- A 2019 guide to Human Pose Estimation with Deep Learning - <https://nanonets.com/blog/human-pose-estimation-2d-guide/>

- In this post, I write about the basics of Human Pose Estimation (2D) and review the literature on this topic. This post will also serve as a tutorial in Human Pose Estimation and can help you learn the basics.
- The MOPED framework: Object recognition and pose estimation for manipulation - <https://personalrobotics.cs.washington.edu/publications/collet2011moped.pdf>
 - We present MOPED, a framework for Multiple Object Pose Estimation and Detection that seamlessly integrates single-image and multi-image object recognition and pose estimation in one optimized, robust, and scalable framework.
 - 2011年 . ディープニューラルネットワーク以前の話 .
- CrowdPose: Efficient Crowded Scenes Pose Estimation and A New Benchmark - <https://arxiv.org/pdf/1812.00324.pdf>
 - In this paper, we propose a novel and efficient method to tackle the problem of pose estimation in the crowd and a new dataset to better evaluate algorithms.
- Pose Detection comparison : wrnchAI vs OpenPose
 - we got a chance to try the state-of-the-art pose-estimation system (wrnchAI) built by wrnch and compare it ' s performance with OpenPose.
- Learning Feature Pyramids for Human Pose Estimation - <https://arxiv.org/pdf/1708.01101.pdf>
 - In this work, we design a Pyramid Residual Module (PRMs) to enhance the invariance in scales of DCNNs. Given input features, the PRMs learn convolutional filters on various scales of input features, which are obtained with different subsampling ratios in a multibranch network.
 - <https://github.com/bearpaw/PyraNet>
- CamLoc: Pedestrian Location Detection from Pose Estimation on Resource-constrained Smart-cameras - <https://arxiv.org/pdf/1812.11209.pdf>
 - In this paper we show that pedestrian location estimation using deep neural networks is achievable on fixed cameras with limited compute resources. Our approach uses pose estimation from key body points detection to extend pedestrian skeleton when whole body not in image (occluded by obstacles or partially outside of frame), which achieves better location estimation performance (inference time and memory footprint) compared to fitting a bounding box over pedestrian and scaling.
- AIIA DNN Benchmark Overview - <https://github.com/AIIABenchmark/AIIA-DNN-benchmark>
 - The goal of the alliance is provide selection reference for application companies, and provide third-party evaluation results for chip companies.
- Performance Analysis of Real-Time DNN Inference on Raspberry Pi - [http://digital.csic.es/bitstream/10261/163973/1/Performance Analysis of Real Time DNN on RPi.pdf](http://digital.csic.es/bitstream/10261/163973/1/Performance%20Analysis%20of%20Real%20Time%20DNN%20on%20RPI.pdf)
 - In this paper, we present a comparative study of some of these frameworks(Caffe, OpenCV, TensorFlow, Caffe2) in terms of power consumption, throughput and precision for some of the most popular Convolutional Neural Networks (CNN) models.
- An update to DeepBench with a focus on deep learning inference - <https://svail.github.io/DeepBench-update/>
 - DeepBench included five basic building blocks of deep learning training: matrix multiply, convolutions, recurrent operations (vanilla and Long Short Term Memory (LSTM)) and all reduce.
- Benchmark Analysis of Representative Deep Neural Network Architectures - <https://arxiv.org/pdf/1810.00736.pdf>
 - This work presents an in-depth analysis of the majority of the deep neural networks (DNNs) proposed in the state of the art for image recognition.
 - <https://github.com/CeLuigi/models-comparison.pytorch>
 - Bianco, Simone & Cadone, Rami & Celona, Luigi & Napoletano, Paolo. (2018). Benchmark Analysis of Representative Deep Neural Network Architectures. IEEE Access. 6. 64270-64277. 10.1109/ACCESS.2018.2877890.
- OpenVINO - <https://software.intel.com/en-us/blogs/2018/05/15/accelerate-computer-vision-from-edge-to-cloud-with-openvino-toolkit>
- Computational complexity of machine learning algorithms - <https://www.thekerneltrip.com/machine-learning/computational-complexity-learning-algorithms/>
- Real-time 2D Multi-Person Pose Estimation on CPU: Lightweight OpenPose - <https://arxiv.org/abs/1811.12004>
 - In this work, we approached the problem of human pose estimation network, suitable for real-time performance on edge devices.
- TBD (Training Benchmark for DNNs) Training Suite

- <https://github.com/tbd-ai/tbd-suite>
- <http://tbd-suite.ai/>
- A 2019 Guide to Object Detection - <https://heartbeat.fritz.ai/a-2019-guide-to-object-detection-9509987954c3>
- A Beginner's Guide to Object Detection - <https://www.datacamp.com/community/tutorials/object-detection-guide>
- DAWN Bench: An End-to-End Deep Learning Benchmark and Competition - <https://dawn.cs.stanford.edu/benchmark/papers/nips17-dawnbench.pdf>
 - <https://dawn.cs.stanford.edu/benchmark/>
 - DAWN Bench is a benchmark suite for end-to-end deep learning training and inference.
- Scalability Comparison Scripts for Deep Learning Frameworks - <https://github.com/awslabs/deeplearning-benchmark>
- Deep Learning Benchmarking Suite - <https://hewlettpackard.github.io/dlcookbook-dlbs/#/>
 - Deep Learning Benchmarking Suite (DLBS) is a collection of command line tools for running consistent and reproducible deep learning benchmark experiments on various hardware/software platforms.
- Benchmarking State-of-the-Art Deep Learning Software Tools - <https://arxiv.org/abs/1608.07249>
- <https://mlperf.org>
 - <https://venturebeat.com/2019/06/24/mlperf-introduces-machine-learning-inference-benchmark-suite/>
- Benchmarking State-of-the-Art Deep Learning Software Tools - <http://dlbench.comp.hkbu.edu.hk/>
- Machine Learning at Facebook: Understanding Inference at the Edge - <https://research.fb.com/wp-content/uploads/2018/12/Machine-Learning-at-Facebook-Understanding-Inference-at-the-Edge.pdf>
 - This paper takes a datadriven approach to present the opportunities and design challenges faced by Facebook in order to enable machine learning inference locally on smartphones and other edge platforms.
- Lecture 6: Modern Object Detection - [https://zsc.github.io/megvii-pku-dl-course/slides/Lecture6\(Object%20Detection\).pdf](https://zsc.github.io/megvii-pku-dl-course/slides/Lecture6(Object%20Detection).pdf)
- Wavenet - <https://modelzoo.co/model/wavenet-tensorflow>
 - This is a TensorFlow implementation of the WaveNet generative neural network architecture for audio generation.
- espnet - <https://modelzoo.co/model/espnet>
 - End-to-End Speech Processing Toolkit [espnet.github.io/espnet](https://github.com/espnet/espnet)
- pytorch-CycleGAN-and-pix2pix - <https://modelzoo.co/model/pytorch-cycleGAN-and-pix2pix>
 - PyTorch implementation for both unpaired and paired image-to-image translation.
- DCGAN-tensorflow - <https://modelzoo.co/model/dcgan-tensorflow>
 - Tensorflow implementation of Deep Convolutional Generative Adversarial Networks which is a stabilize Generative Adversarial Networks.
- Open-source (MIT) Neural Machine Translation (NMT) System - <https://modelzoo.co/model/open-source-mit-neural-machine-translation-nmt>
 - This is a Pytorch port of OpenNMT, an open-source (MIT) neural machine translation system. It is designed to be research friendly to try out new ideas in translation, summary, image-to-text, morphology, and many other domains.
- Chatbot - <https://modelzoo.co/model/chatbot>
 - Implementation of "A neural conversational model"
- Jetson Nano: Deep Learning Inference Benchmarks - <https://developer.nvidia.com/embedded/jetson-nano-dl-inference-benchmarks>
 - model, Jetson Nano, RPi3, RPi3+Movidius2, TPU Dev Board
 - ResNet-50(224x224), 36 FPS, 1.4 FPS, 16 FPS, DNR
 - MobileNet-v2(300x300), 64 FPS, 2.5 FPS, 30 FPS, 130 FPS
 - SSD ResNet-18(960x544), 5 FPS, DNR, DNR, DNR
 - SSD ResNet-18(480x272), 16 FPS, DNR, DNR, DNR
 - SSD ResNet-18(300x300), 18 FPS, DNR, DNR, DNR
 - SSD Mobilenet-V2(960x544), 8 FPS, DNR, 1.8 FPS, DNR
 - SSD Mobilenet-V2(480x272), 27 FPS, DNR, 7 FPS, DNR
 - SSD Mobilenet-V2(300 × 300), 39 FPS, 1 FPS, 11 FPS, 48 FPS
 - Inception V4(299 × 299), 11 FPS, DNR, DNR, 9 FPS

- Tiny YOLO V3(416 × 416), 25 FPS, 0.5 FPS, DNR, DNR
- OpenPose(256 × 256), 14 FPS, DNR, 5 FPS, DNR
- VGG-19(224 × 224), 10 FPS, 0.5 FPS, 5 FPS, DNR
- Super Resolution(481 × 321), 15 FPS, DNR, 0.6 FPS, DNR
- Unet(1x512x512), 18 FPS, DNR, 5 FPS, DNR
- ReForm: Static and Dynamic Resource-Aware DNN Reconfiguration Framework for Mobile Device - <http://mason.gmu.edu/~lzhao9/materials/papers/a183-Xu.pdf>
 - we propose ReForm - a resource-aware DNN optimization framework. Through thorough mobile DNN computing analysis and innovative model reconfiguration schemes (i.e. ADMM based static model fine-tuning, dynamically selective computing), ReForm can efficiently and effectively reconfigure a pre-trained DNN model for practical mobile deployment with regards to various static and dynamic computation resource constraints.
- Edge Intelligence: On-Demand Deep Learning Model Co-Inference with Device-Edge Synergy - <https://arxiv.org/pdf/1806.07840.pdf>
 - we propose Edgent, a collaborative and on-demand DNN co-inference framework with device-edge synergy.
- Learning to Predict Depth on the Pixel 3 Phones - <https://ai.googleblog.com/2018/11/learning-to-predict-depth-on-pixel-3.html>
- ニューラル・ネットワークにおける情報量の表現と FPGA 実装 - <https://www.ipsj-tokai.jp/jigyoku/files/H29slide20171124.pdf>
- Machine Learning on FPGAs - http://cadlab.cs.ucla.edu/~cong/slides/HALO15_keynote.pdf
 - Feedforward computation on FPGA
- A Survey of FPGA-Based Neural Network Inference Accelerator - <https://dl.acm.org/citation.cfm?id=3289185>
 - A guide to AI-Chip-related articles (on my Wechat blog) - <https://medium.com/@shan.tang.g/a-guide-to-ai-chip-related-articles-on-my-weichat-blog-5ffc440950d3>
- AI Chip (ICs and IPs) - <https://github.com/basicmi/AI-Chip>
- AI on FPGAs - https://www.omnitek.tv/dpu_fpga
- Deep Learning Hardware: FPGA Vs. GPU - <https://semiengineering.com/deep-learning-hardware-fpga-vs-gpu/>
- On-Device Neural Net Inference with Mobile GPUs, <https://arxiv.org/pdf/1907.01989.pdf>
- MAIX performance and limit for AI models. - <https://bbs.sipeed.com/t/topic/691>