# FINLAY HUDSON

**Lead Computer Vision Researcher & Engineer with a strong track record in solving customer problems by building upon state-of-the-art research to bring bespoke solutions to production. Proficient in Agile project planning & ceremonies, ML Engineering in Python, mentoring Junior Engineers & Researchers as well as drafting Statements of Work & Proposals for AI consultancies. Committed to pushing the boundaries of computer vision for exciting real-world use cases.**

## CONTACT


finlay.hudson@hotmail.com
07831713291
Hitchin, Hertfordshire
Finlay Hudson
finlay-hudson.github.io


## SKILLS

*Programming*

**Python**
**C++**
**CUDA**
**SQL**
**MATLAB/LUA**
**R**

*Operating Systems*

**Linux**
**MacOS**
**Windows**

*Software, Libraries & Tools*

**Visualisation**
(matplotlib, plotly, dash, imgui, Tkinter, TensorBoard, DataDog)
**Data Manipulation**
(numpy, scipy, pandas, cupy, numba)
**Image Processing**
(opencv, pillow, scikit-image)
**Docker**
**Git / GitLab / Bitbucket**
**Cloud Services**
(AWS (Kinesis/S3/EC2), Google Cloud Platform)
**CAD**
(AutoCAD, Inventor)
**EAGLE**
**LaTeX**

*Deep Learning Architectures*

**2D/3D Segmentation**
(SAM2, Deeplab, Mask RCNN, PSPNet)
**Object Detection**
(YOLO, DETR, Fast RCNN, HydraNets)
**GAN**
(Pix2Pix, InPainting, CycleGAN, NeRF)
**Transformers**
(ViT, Swin)
**Action Recognition**
(TimeSFormer, MoViNet)
**LLM/NLP**
(LLaMA, GPT, CLIP*text-image*, BERT)
**Diffusion**
(Stable Diffusion, Video Stable Diffusion, COCOCO)

*Data Analyzing / Annotation*

**V7 Darwin**
**FiftyOne**
**Neptune**
**Weights and Biases**

## WORK HISTORY

July 2023 - Present
Spectral Compute Ltd.
**GPU Software Engineer (Part-time)**

*Implementing deep learning model optimisation strategies as well as supporting the building of internal libraries.*

- **Model Optimisation** – Using ONNX, TensorRT, TensorRT-LLM as well as custom CUDA modifications I have optimised clients deep learning models to achieve fast inference times as well as keeping good results and also allowing for flexibility in terms of settings, input sizes, batch sizes etc. This has been done on SDXL, LLaVA and CogVLM pipelines. Speeding up pipeline inference compared to default implementation speed by up to 40%.

---

February 2024 - Present
Owl + Lark Ltd.
**Lead AI Engineer (Part-time)**

*Using cutting-edge innovations to design a product offering holistic sleep analysis and tailored recommendations for optimal sleep hygiene.*

- **Non-invasive Sleep Analysis** - Designed and implemented a system for radar-based detection of vital signs and actimetry data, leveraging this data to train a machine learning model for predicting individual sleep stages throughout the night.
- **Holistic Sleep Recommendations** - Developed a custom LLaMA-based LLM system to analyze user data—including circadian rhythms, daily activities, sleep habits, vital signs, and actimetry metrics—to provide tailored recommendations to improve sleep quality.

---

June 2022 - July 2023
Satis AI Ltd.
**Lead Computer Vision Engineer**

*Lead of a small team tasked with creating a real-time, human in the loop, visual analysis system working within large-scale fast-food kitchens, including those of a Fortune 500 company. Responsible for solving technical challenges within the system; improving accuracy, speed and reducing resource costs, as well as conducting project management and mentoring graduate/junior researchers. Resulting in an robust perception system capable of running on the cloud or edge devices with a p95 latency of <100ms.*

- **Action Recognition** – Led a team of 5 to develop an action recognition model for labor management; orchestrating planning sessions, conducting code reviews, and assigning tasks. Aligned system requirements with project objectives in collaboration with the product manager. Leveraged a pre-trained TimeSFormer model to train a basic classifier, saving time and money by avoiding custom dataset creation. Integrated the model into the pipeline using Multiprocessing without compromising system latency. Achieved 95% accuracy for cleaning tasks, further improved with temporal smoothing techniques.
- **Kitchen Tracker** - Developed a real-time multi-object tracker for fast-food kitchens capable of tracking humans, ingredients and packaging. Optimised performance with efficient vectorised Python code, achieving a 3-5x speed improvement. Successfully deployed across 15 unique kitchen locations. By leveraging a wide array of contextual data within the scene, the tracker accurately identifies objects over time, enabling efficient monitoring and management in fast-food environments.
- **Human Pose Estimation** – Customised and productionised a top-down DarkPose model. By customising the number of pose keypoints, I enabled full hand detection, leveraging the estimated heatmaps to achieve fast and high-quality hand segmentation. Human data was sampled with strategies to ensure uniqueness and variation of data before leading an annotation team on annotation strategy.
- **Model Inference Optimisation** – Optimized deep learning models using TensorRT FP-16, reducing latency by 66% and costs by 50%. Developed GPU-based routines and deployed Triton Ensemble models, improving throughput and cutting resource costs by 30%. Integrated DALI for efficient preprocessing, achieving a 70% CapEx reduction for clients.
- **Data Sampling Framework** – Developed a framework to gather feedback from later pipeline stages for sample images where object detection performed poorly. Annotators reviewed these images for model retraining in an online learning fashion, ensuring continuous adaptation to dynamic kitchen environments.
- **Custom Model Creation** – Created a custom model architecture and training pipeline to allow for a single feature extractor backbone to be used for object detection, instance segmentation and human pose estimation.
- **Product Orientated Metrics** – Created custom metrics to allow for product driven development throughout a system when "standard" research orientated metrics would not be sufficient. These metrics were further expanded to establish key performance indicators (KPIs) at the product level, enabling a holistic evaluation of system performance aligned with product objectives.

*Model Optimisation*

**TensorRT Inference**
**ONNX/ONNX Runtime**
**TorchScript**
**Nvidia Ecosystem**
(Triton, Ensemble, DALI, CUDA)

*Hardware*

**Dynamixel**
**PixHawk**
**Sensors**
(Time of Flight, LiDAR, Radar ...)
**Intel Real Sense**

*CI/CD and Provisioning*

**Unit Testing**
**Python Packaging**
**Concourse CI**
**GitHub Actions**
**Rake**
**Terraform**

*Management and Documentation*

**Management Tools**
(Jira, Asana, Click-up)
**Miro**
**Documetation**
(Notion, Confluence)

## LOOKING FOR

I am interested in Deep Learning based roles with a focus on Computer Vision, open to full time and part time contracts as well as consulting work.

## REFERENCES

Available on request.

August 2021 - June 2022
Springbok AI Ltd.

**Lead Computer Vision Engineer**

*Delivered high quality bespoke solutions for both computer vision and NLP based clients. Leading the production all AI components.*

- **Sentiment Analysis** - Used BERT to gauge the sentiment towards a variety of different chatbot experiences. Scraping data from a large variety of uncurated online sources.
- **Chatbot Creation** - Utilised GPT-3 to generate components required in creating a chatbot using RASA.
- **AI Assisted Sports Analysis** - Developed technology for automated race reporting, focusing on jockey identification and position understanding. For jockey ID, implemented a pipeline to detect and identify numbers on saddle-cloths and used a Siamese network with contrastive loss for accurate matching of jockey silks. For position understanding, trained a Detectron network and customized SORT tracker for temporal tracking, enhanced with fine-tuned Kalman filtering. Employed movement perspective calculations to determine horse ranking and integrated scene change detection for tracker re-calibration.

January 2020 – August 2021
Visio Impulse Ltd.

**Computer Vision Scientist**

*Lead developer of projects while part of a small team to solve a variety of deep learning problems for a range of clients. Mostly computer vision based problems as well as some NLP tasks.*

- **Underground Component Analysis** - Developed an AI-driven solution for a multinational energy company to verify underground component placement on construction sites. Using deep-learning-based semantic segmentation, components were detected in geo-referenced drone ortho-images and compared to as-built and issued-for-construction CAD drawings. Applied time-series analysis to select reliable geospatial timestamps, enabling accurate autonomous deviation assessments.
- **AI Assisted Satellite Imagery Tagging** – Created a Tkinter platform for efficient semantic segmentation labeling for a large UK mapping agency. A PyTorch based deep learning back-end is there to aid tagging and increase speed. Due to the data being homogeneous, once some tagging is manually done the AI back-end learns many of the class definitions, which resulted in much quicker tagging for the user. Plus an already trained model is available if desired. The AI back-end embeds consistency in the tagging.
- **Visual Search** – Developed the back-end for a platform which allowed for a user to select an area of a large scale geo-tagged image, before a trained model then returned areas with similar features to the chosen area. Use case was for quick inventory and safety in uncontrolled environments.

November 2018 – January 2020
HayBeeSee Ltd.

**Robotics Engineer**

*Worked in a small team tasked with designing and building a drone to undertake agricultural weed killing and statistical analysis crucial for farmers.*

August 2018 – November 2018
University of York

**Robotics Research Software Engineer**

*Part of the 'University Research Pump-Priming Fund - RoboSpartan' project. Created a statistical analysis app using Shiny and R, comparing real-life results with the ARGoS simulation results.*

## RESEARCH

*2021 - Present*
University Of York, York

**PhD, Scene Understanding with Zero Shot Learning**

My research focuses on computer vision-based scene understanding, emphasising the integration of non-visual features and semantic or temporal cues. I aim to minimise reliance on task-specific data, reducing the need for extensive dataset curation, to enhance the flexibility and adaptability of computer vision systems for diverse applications. I am particularly interested in test-time training, large foundation models, and diffusion-based techniques, primarily for addressing segmentation and tracking tasks in video data. For publications see: finlay-hudson.github.io

## EDUCATION

2017 - 2018
University Of York, York

**MSc, Intelligent Robotics; Distinction**

Swarm Intelligence - 70%, C programming - 85%, Control Systems for Robotics (Double Credit) - 68%, Practical Robotics (Double Credit) - 70%, Neural Networks for MSc - 67%, Critical Literature Review (Triple Credit) - 74%, Final project (Quadruple Credit) - 77%

**Project**: A group project with the goal of developing a robotic system, incorporating a drone and ground robot, to autonomously complete a search and rescue task.

- Used ROS packages to allow for full control of the drone to be available from a user's computer and keyboard.
- Integrating OpenCV with ROS. Using OpenCV to allow for the drone's camera to detect objects of a certain colour, then, via blob detection, calculate the angle and distance required to allow a gripper to grab an object.
- Used localisation of objects combined with object orientation; the relationship between them was mathematically established.
- Autonomously flight developed, including controllers to reduce drift.

2014 - 2017
Anglia Ruskin University, Cambridge

**BSc (Hons), Audio and Music Technology; 1st Class + Richer Sounds Award**

Mathematics for Technology - 85%, Analog Electronics - 82%, Electronic Circuits - 67%, Live Sound Engineering - 82%, Audio Electronics - 69%, Final Project - 77%, Audio Programming - 80%

**Project**: Incorporated a low-frequency oscillator (LFO) feeding a voltage-controlled filter (VCF) into a pedal for use with an electric guitar.

- Designed and developed an electronic circuit from scratch, using Multisim to prototype the system.
- Allowed the circuit to have different waveform settings to allow the user to create their desired sound.
- Ensured quality wiring and components were in the system to reduce unwanted noise.