

Joshua Thorpe

FULL-STACK AI/ML · CLOUD ARCHITECT · PROJECT MANAGER

Singapore, Singapore

✉ thorpe.josh@outlook.com | 🏠 joshua.thorpe.work | 🗣️ thorpejosh | 📄 joshua-thorpe

Skills

AI/ML	PyTorch, Keras, TensorFlow, Scikit-learn, pandas, NLTK, FastAPI
MLOps	Docker, Kubernetes, Ansible, Jenkins, Postgres, Redis, Mongo, SQLServer, Hadoop, TensorFlow TFX, AutoML
Programming	Python, Go, C++, Bash, LaTeX, YAML/JSON/TOML
Dev Envs	Emacs (Evil), Vim, JupyterLab, PyCharm, VSCode
Hardware/OS	RPI, Arduino, Arch Linux, Debian, RHEL/Fedora/Amazon, openSUSE, WSL, MacOS
Project Mgmt	Jira, Azure Devops Boards, GitHub Projects, Trello, OnePlan, Emacs OrgMode
Cloud	Azure, AWS & GCP: VPS Compute, Kubernetes, Container Instances, Serverless Compute, Cloud Storage, Lake House services, AutoML, IAM, Virtual Networks, VPN, Gateways, Firewall, API & Queue Services, Monitoring/Alarms and Notification services

Education

RMIT (Royal Melbourne Institute of Technology)

Melbourne, Australia

BENG (HONOURS) IN ADVANCED MANUFACTURING AND MECHATRONICS

- First Class Honours — GPA of 3.9/4.0

Honors & Awards

2019	Vice-Chancellors List Award , for graduating in top 2% of RMIT higher education students in 2019	RMIT, Australia
2018	Engineers Australia Prize , awarded to the top engineering student in their final year	RMIT, Australia
2018	IET Student Prize , awarded for Capstone project innovation and excellence	RMIT, Australia
2017	Interpack Advanced Manufacturing Prize , for best Mechatronics student in 2017	RMIT, Australia
2015	Golden Key International Honours Society Membership , for being ranked in top 10% of cohort	International
2014	Australian Tertiary Admission Rank of 97.7 , Achieving top 2.3% of Australian cohort	Vic, Australia
2013	State Premiers award , for a perfect score in VCE Systems Engineering	Vic, Australia

Certifications

2024	AWS Certified Machine Learning — Specialty , Amazon Web Services (AWS) 🔗
2023	AWS Certified Solutions Architect — Associate , Amazon Web Services (AWS) 🔗
2023	AWS Certified Cloud Practitioner , Amazon Web Services (AWS) 🔗
2022	Scrum Master Certification , LearnQuest 🔗
2022	Google Project Management — Professional Certificate , Google 🔗
2022	MLOps — Machine Learning Engineering for Production , DeepLearning.AI 🔗
2019	LE-1 Linux Essentials , Linux Professional Institute 🔗

Projects

Home Lab & Self Hosted Servers

Australia/Singapore

SOLE-OWNER/ ENGINEER

Jun. 2012 – Present

Project goal: Establish a set of servers used for self-hosting storage, cloud services, CI/CD, development environments, various client projects & more.

- 8 bare-metal servers and 1 cloud VPS across 4 locations, 15 users.
- Running Debian, Arch, Fedora, RHEL and FreeBSD.
- Servers networked with WireGuard VPN, OPNsense firewalls with Intrusion detection/prevention, self-hosted internal DNS & ad-blocker.
- Ingress to hosted services via Cloudflare, Traefik reverse proxies, Automated LetsEncrypt SSL certificates, and Authelia MFA authentication.
- Static websites hosted on AWS S3 behind cloudfront CDN, with R53 DNS failover to bare-metal servers.
- Prometheus network and server resource monitoring and metric collection, and Grafana visual dashboards of time-series data.
- Jenkins CI/CD Pipeline for automation of private and open-source projects.
- 50+ Docker and Kubernetes containers running across multiple machines.
- 20+ custom container applications to handle automation, alerts, backups. Several Signal/Telegram bots for easy customer interfacing with private APIs.
- Remote video rendering, GPU compute and deep learning services.
- Managed/deployed with Ansible automation.
- Regularly perform penetration testing/ ethical hacking to ensure security.

MLOPs Project for Siri Language Engineers

Singapore

MLOPS & DEVOPS ENGINEERING

Jun. 2022 – Present

Project goal: Accelerate Language Engineers development time by providing automated feedback on code changes.

- Built containerised tooling to provide feedback on syntax, formatting, duplication, etc.
- Integrated tooling into CI/CD pipeline to provide automated feedback on code changes and the predictive influence on the language model through a summary of changes in the training dataset.
- Developed a bash helper script that ran the containerised pipeline locally on macOS, greatly accelerating development for Language Engineers.
- Helper script managed git for Engineers, calculating the diffs between local working branch and upstream main, running only the applicable tests from the pipeline, saving up to 2 hours per code change.
- Achieved 100% adoption rate of the local tooling from Language Engineers in the Global Siri team due to the increased rate of feedback and agility it granted them.

Software Development Platform, CI/CD, LakeHouse and Data Ingestion Pipeline.

Sydney, Australia

MLOPS & AND PLATFORM ENGINEERING

May. 2021 – Jun. 2022

Project goal: Architect, develop and deploy a complete software development platform, and big data ingestion & storage platform to facilitate the startup of multiple software projects within a medical device company.

- Defined new software project Agile standards and best practices, overhauling existing code to comply.
- Setup a software development platform including cloud environment, CI/CD pipelines, project planning and collaboration software, documentation & training.
- Setup a LakeHouse cloud environment, ingesting data from on-prem manufacturing IOT devices and new orders submitted from global hubs.

3D Dental Scan Anomaly Detection

Sydney, Australia

ML & COMPUTER VISION

May. 2021 – Jun. 2022

Project goal: Detect issues in 3D dental scans on patients teeth to prevent manufacturing of sub-optimal dental devices.

- Modelled and trained neural networks to automatically detect if 3D dental scan is from upper or lower jaw, then segment and classify each tooth individually, allowing automated decisions on product compatibility for each customer.
- Developed machine learning algorithms to detect if 3D dental scan is incomplete, corrupted, or contains artifacts such as foreign matter that render it unusable for manufacturing a dental appliance.
- Lead and developed algorithms that simulate 3D jaw movements, dental bite positions and detects problematic dental scan geometries prior to manufacturing a customer device.

Research and Analysis of Dental Appliance Forces on Teeth

Sydney, Australia

COMPUTER VISION & DATA SCIENCE

May. 2021 – Jun. 2022

Project goal: Research and analyse what factors contribute to the force required to remove a dental appliance from a patients teeth, and develop a strategy for controlling this force for reliable manufacturing of safe & comfortable oral appliances.

- Managed a deep learning approach that used the distances between 3D models of teeth and appliances to infer high and low force regions.
- Successfully proved the companies first deep learning approach on 3D meshes with the model accurately learning regions of interest on models of oral appliances.
- Analysed millions of data points with data science toolkits and created an interactive report, visualising the findings with heat-maps and simulations. Report was built with Jupyter Books behind an nginx container and hosted on Azure.


Automated Timelapse Generator

Melbourne, Australia

COMPUTER VISION ENGINEER

Sep. 2018 – Sep. 2019

Project goal: Develop software to automate the generation of timelapse videos given a set of raw pictures taken from long-term timelapse cameras.

- Developed algorithms to detect and remove images that were blurry, over/under exposed, had dirt on the lens, etc.
- Engineered a stabilisation algorithm that aligned images using common points & homogeneous transformations, to account for camera movement between photos causing image misalignment.
- Built a blending filter that blended images across a 1 second window to smooth out the video image and avoid flickering from intermittent objects.
- Manipulated image color, exposure, and gamma to match the rolling average of the surrounding images in the video sequence to provide a smooth and professional timelapse video.
- Built and deployed time-lapse generator on scalable cloud infrastructure using AWS EC2, S3, SQS queues and custom autoscaling policies.
- The released automated time-lapse generator  was earning the small business over \$50K of revenue per month.

Automated Safety Reports for Construction Sites

Melbourne, Australia

COMPUTER VISION ENGINEER

Sep. 2018 – Sep. 2019

Project goal: Use long-term timelapse camera images to detect and report safety issues within construction sites to the Project Manager.

- Retrained RCNN model to detect and classify safety related objects such as people, safety vests, gloves, hard hats, boots, etc and machinery related objects such as scissor lifts, cherry pickers, backhoes, concrete mixers, excavators, etc.
- Developed an alert software that monitored the object detection results and alerted based on rule-sets such as; person on a scissor lift without a hard hat.
- Used exploratory data analysis to identify metrics from the object detection data such as worksite start/stop time, total man hours worked, and productivity metrics based on behavioural analysis.
- Generated automated visual reports to summarise and track the data relating to safety and custom metrics extracted from the data. Allowed monitoring trends on safety, attendance, productivity and progress on the construction site.
- Report included machinery metrics such as which types and count of machinery were present on which days, allowing easy inventory/order validation.

Medical Cellular Image Segmentation for Cancer Detection

Melbourne, Australia

HONOURS CAPSTONE PROJECT

Mar. 2018 – Nov. 2018

Project goal: Develop a robust deep learning method to solve the 60 year old problem of autonomously segmenting microscopic images of cells.

- Successfully developed a pixel-wise segmentation CNN trained on only 30 images, achieving an accuracy of 98.2% on 2000 test images.
- Identified a unique property of ovarian cells mitosis through exploratory data analysis that helps to track and identify cancerous cells.

Industry 4.0 Meat Processing

Melbourne, Australia

ECP SUMMER INTERNSHIP — INTELLIGENT AUTOMATION RESEARCH GROUP

Nov. 2017 – Feb. 2018

Project goal: create an Industry 4.0 autonomous system for cattle washing and lameness detection for the Australian Meat Processing Industry.

- Developed software to capture and synchronise data from 4 Kinect sensors, 2 thermal flares and 4 capacitive pressure sensors simultaneously.
- Aligned multiple point clouds to generate a 3D model of cattle with colour projection and thermal information.

Employment History

DBS Bank

Singapore

MACHINE LEARNING ENGINEER — GLOBAL FINANCIAL MARKETS

May. 2024 – present

- Researched and developed a dynamic FX pricing algorithm, leveraging EDA and A/B testing, resulting in an estimated SGD\$5M annual profit increase.
- Developed and deployed a GenAI SOP automation tool, estimated to save the Operations team the equivalent of 20 FTEs.
- Architected and built microservices for both projects, integrating diverse data sources and ensuring production readiness.
- Established a robust evaluation and monitoring frameworks; utilising data drift, judge LLMs and user feedback for continuous improvement.
- Standardised team coding practices and CI/CD pipelines using Jenkins, Package Managers, Testing, Linting/Formatting and Vulnerability Scanning, improving code quality and accelerating deployment.

Apple

Singapore

MACHINE LEARNING ENGINEER — GLOBAL SIRI

Jun. 2022 – Jun. 2023

- Responsible for end-to-end Siri model performance for Australian and New Zealand (ANZ) markets.
- Identified and spearheaded business and feature opportunities, including improving indigenous & minority language support for ANZ, Canada, India, Singapore and South Africa.
- Added ANZ support for Top 200 most requested emojis, automated grammar correction, and new iOS integration features.
- Significantly reduced ANZ issue count by improving Maori Language support.
- Reviewed and curated recruitment material, and helped with on-boarding and mentoring junior machine learning engineers.

SomnoMed

Sydney, Australia

AI/ML PROJECT ENGINEER — RESEARCH & DEVELOPMENT

May. 2021 – Jun. 2022

- Managed multiple computer vision projects (budget totalling \$400k) from initiation to execution.
- Worked with key stakeholders to define project requirements and seek regular feedback for responsive improvements.
- Defined new software project Agile standards and best practices, overhauling existing projects to comply.

RMIT University

Melbourne, Australia

RESEARCH ASSISTANT — INTELLIGENT AUTOMATION RESEARCH GROUP

Jan. 2020 – Dec. 2020

- Researched and identified rotation averaging gaps in the field of computer vision with applications such as: crime scene recreation; real-time multi-camera calibration on autonomous vehicles or manufacturing/sorting facilities; or recreating 3D models of scenes from 2D pictures.
- Developed, tested, and analysed new state-of-the-art graph deep learning methods using frameworks such as PyTorch and TensorFlow.
- Wrote an academic paper summarising the method, its benefits and how it compares to existing methods.
- Developed a novel method of performing rotation averaging using deep learning, that is more accurate, more robust and orders of magnitude faster than previous methods.
- Results unlocked ability for large 3D scenes to be recreated from hundreds to hundreds of thousands of 2D images in real time.

Daimler AG (Mercedes Benz)

Stuttgart, Germany

INTERN CONTROL ENGINEER — AUTONOMOUS DRIVING RESEARCH & DEVELOPMENT

Jul. 2019 – Dec. 2019

- Achieved an accurate data-driven LTI model of the Mercedes S-Class drive-train through system identification and machine learning, greatly expanding the possibilities for position controllers.
- Successfully developed a LQR position controller for the Mercedes S-Class autonomous vehicle, that performed comparably to the existing controller and saved the research team significant amounts of time due to its greater tunability.

Cornerstone Solutions

Melbourne, Australia

COMPUTER VISION ENGINEER — PRODUCT DEVELOPMENT TEAM

Sep. 2018 – Sep. 2019

Macedon Ranges Shire Council

Macedon Ranges, Australia

OUTDOOR POOL MANAGER

2015 – 2018

Extracurricular Activity

Contributing to Open-Source Projects

International

CONTRIBUTOR AND/OR OWNER

2018 – Present

Official Raspberry Pi Python library · Containerised Linux Borg backup & web UI · Multi-threaded media sorting tool · IOT smart automated motorised gates · Arduino tunable turbo engine management software · Deep Declarative Networks

RMIT Formula SAE Team

Melbourne, Australia

ENGINEER ON POWERTRAIN AND ELECTRONICS TEAM

2018 – 2019

- Engine management system design and wiring.
- Engine testing and dyno tuning.

Humanitarian Engineering

Tanzania, Cambodia, Nepal

HUMANITARIAN ENGINEER

2011, 2016, 2018

- Nepal: Developed a powerful transport system for safely moving building supplies across remote mountainous terrain, using readily available materials. Helped setup a local workshop to start producing the system.
- Cambodia: Engineers Without Borders design summit, Designed an appropriate waste disposal mechanism in a remote village on that prevented further contamination to their water supply.
- Tanzania: Designed safer electrical infrastructure and educated children and staff at several schools and orphanages on electrical safety.

RMIT LEAD Academic Mentoring — 3rd Year Automatic Control Systems

Melbourne, Australia

MENTOR & TUTOR

2018

Toastmasters Sunbury and Macedon Ranges Club

Melbourne, Australia

MEMBER

2015 – 2016


Bendigo Yacht Club

Bendigo, Australia

VOLUNTEER INSTRUCTOR

2011 – 2016

Interests

 Linux · Contributing to open-source projects · Expanding home lab · Ethical hacking · Building custom ergonomic keyboards · Fitness