Angad Ahuja

github.com/AngadSingh22

Email: angad.ahuja_ug2023@ashoka.edu.in

Mobile: +91 9896866733

Education

- **Ashoka University** — Sonipat, Haryana, India
  B.Sc. Computer Science & Mathematics — *2023 – 2028*

Experience

- **Centre for Constrained Image Generation** — Columbus, Ohio, USA
  Post-Generation Team — *Oct 2025 – Present*
  - **Constraint-aware generation pipelines**: Built constrained image-generation workflows (diffusion and inpainting with deterministic slot-filling), developed and evaluated 6 architectures end-to-end, and selected a final pipeline that reliably produced 600 DPI and 300 DPI outputs as required with compliance checks and export-validity gates.
  - **Evaluation and reliability**: Designed an evaluation harness tracking fidelity metrics such as SSIM and PSNR where applicable, and logging inference latency and throughput, to guide iteration toward stable deployment-ready systems.
  - **Team formation**: Created and led a team of 11, establishing sprint structure, code-review gates, and reproducible experiment tracking to increase delivery velocity and reduce regressions.
- **Science of Sadharan with Prakash Labs (Stanford)** — India
  Research Lead, M&E Vertical — *Jul 2024 – Oct 2025*
  - **NLP for field feedback**: Spearheaded NLP pipelines (custom tokenization + Transformer-based sentiment models) to convert field feedback into analysis-ready signals for reporting workflows.
  - **Vision models for documentation**: Designed and trained CNN-based image-labelling models, achieving 92% validation accuracy, to classify educational materials and on-site demonstration photos.
  - **Data collection automation**: Architected and deployed a WhatsApp chatbot (Twilio API + Python) to standardize real-time remote data collection with Firebase-backed storage and state management, collected data from over 2000 students.
  - **Multilingual analysis**: Conducted Hindi/English normalization and topic modelling (LDA + Word2Vec) to derive interpretable themes supporting impact evaluation.
  - **Operations and reporting**: Managed end-to-end M&E workflows for a five-member team, conducted qualitative coding in NVivo, and published comprehensive Impact Reports following each volunteer series.
- **KCDH-IITB (Koita Centre for Digital Health)** — Mumbai, India
  Research Intern — *October 2023 – Jul 2024*
  - **Mental-health support chatbot**: Built a pipeline processing PHQ-9 and GAD-7 responses through a Transformer-based sentiment classifier to generate affective-state vectors, and iterated using engagement and response-fidelity checks in IRB-approved pilots.
  - **Agent-based computational model**: Developed an agent-based simulation to evaluate policing, screening, and therapeutic interventions for university substance use, validated through sensitivity analyses to support harm-reduction recommendations, and formed a team of 5 to execute model development and evaluation.

Projects

- **Deepfake Detection on Edge (ResNet-50 + Transfer Learning + Quantization) (2025)**: Built an edge-deployable deepfake image detector using ResNet-50 transfer learning, then quantized for deployment, achieving 82% accuracy on the quantized model while tracking latency and size trade-offs to keep inference practical on constrained devices.
- **ROI Color Control Losses for Inpainting (Inference-Time) (2025)**: Implemented severe colour-control inference-time loss functions, and monitored per-step ROI drift and reconstruction error to preserve target color fidelity while keeping the background unchanged, achieving ΔE_{00} under 5.
- **Development of a Weighted Transparency Index for Indian Non-profits (with CSIP, Ashoka) (Dec 2024 – May 2025)**: Synthesised mandatory and voluntary reporting dimensions using fairness-constrained optimisation to derive weightings penalising asymmetrical disclosures, collected over 68 transparency items, verified scores against Guidestar India rankings, and led a team of 3 in collaboration with CSIP at Ashoka.
- **Vision Transformer (ViT) on CIFAR-10 (Jun 2025)**: Built a ViT from scratch with patch embeddings, attaining 75% test accuracy with a 12% improvement over a CNN baseline, reduced training time by 22% via gradient checkpointing, and implemented LR warmup with cosine decay to converge 3x faster than standard schedulers.

Skills

- **Languages**: Python, C++, R, MATLAB
- **ML Frameworks**: PyTorch, TensorFlow, Hugging Face Transformers
- **Edge AI / Deployment**: TensorFlow Lite (TFLite), PTQ/QAT quantization, pruning, knowledge distillation, Raspberry Pi-class deployments
- **Tools**: Git, OpenCV, spaCy, FlashAttention
- **Domains**: Deep Learning (NLP), Computer Vision, Diffusion/Inpainting, CNNs, Vision Transformers, Edge Inference, Formal Verification Methods