

# Koustav Ghosal

Dublin, Ireland | +353 892359796 | [mail.koustavghosal@gmail.com](mailto:mail.koustavghosal@gmail.com)  
[Webpage](#) | [GitHub](#) | [Google Scholar](#)

---

## Bio

---

I am a Lead AI/ML Engineer (Manager) at Mastercard. In my current role, I am part of an AI product enablement team that transitions ideas from R&D into scalable stage deployments and supports development teams in production rollout. Additionally, I lead applied R&D initiatives in graph machine learning, agentic computing, and foundation models within the financial services domain.

Prior to Mastercard, I worked at Accenture as a research scientist. I hold a Ph.D. in Computer Science from Trinity College Dublin and a research masters degree from IIIT Hyderabad, India. I have about 5 years of industry and 7 years of academic experience in machine learning.

---

## Experience

---

### **Lead AI/ML Engineer (Manager) | Mastercard, Dublin | December 2024 - Present**

- Scale and evaluate research prototypes in staging on near real-time data; collaborate with data scientists to monitor drift using observability tools like Fiddler and Grafana.
- Build and maintain on-prem and cloud-based data pipelines for legacy (Hadoop) and streaming (Kafka) data across multiple availability zones in coordination with data engineers.
- Design and review production-grade software architectures; deploy Kubernetes-based AI models for security use cases, ensuring compliance, resiliency, and scalability in financial services.
- Lead R&D in Generative AI with focus on Graph ML, Agentic Computing, and foundation models; mentor teams, track emerging trends, and drive publications and patents. Collaborate with internal and external partners such as data-science teams, universities for R&D initiatives.

---

### **Research Scientist (Assoc. Manager) | Accenture, Dublin | January 2022 - December 2024**

- Led MLOps development in a large cross-functional team; architected a multi-agent system for business analytics to support scalable, automated decision-making.
- Developed healthcare prototypes leveraging structured and unstructured EHR data for predictive modeling and clinical workflow automation.
- Mentored internal teams on AI engineering practices and engaged with clients to align solutions with emerging AI capabilities.
- Co-authored a research paper on AI-driven business analytics and filed a patent for a medical imaging AI tool.

---

### **Postdoctoral Researcher | V-SENSE, Dublin | March - December 2021**

Explored vision transformers and graph neural networks for 3D human pose estimation and 3D mesh sequence compression for augmented and mixed reality applications.

---

### **Research Engineer | Deep Learn Labs, India | August, 2015 - March, 2016**

Developed a deep learning-based prototype for vehicle detection and classification using OpenCV, Caffe, and C++.

---

### **Software Developer | OMitra, India | August 2014 - January 2015**

Developed live in-journey features for the proof of concept such as chats, meal and cab bookings (Android utility app for train journeys)

---

## Education

---

### PhD in Computer Science

Trinity College Dublin, Ireland, 2017-2021

During my PhD, I worked on Computer Vision and NLP for exploring the capacity of AI to analyze photographic images. Using a large, real-world dataset of images and user comments, I built tools combining CNN, LSTM and Graph Neural Networks for applications such as image attribute prediction, image captioning, and aesthetic score regression.

Previously, I did a Research Master's in Computer Vision from IIIT-Hyderabad, India (2016) and a B-Tech in Computer Science from West Bengal University of Technology, India (2011).

---

## Skills

---

**Subject Areas:** Computer Vision, Natural Language Processing (NLP), Graph Machine Learning, Large Language Models (LLM), Optimisation methods, Classical ML (such as SVMs, Random Forests, Bayesian Methods), Digital Signal Processing, Database Systems, Agentic AI, RAG, Microservices, Cloud Computing, Software Engineering.

**Programming:** Python, C, C++, Java, Matlab, PyTorch, Tensorflow, GitHub, Docker, Kubernetes, Spark, Kafka, Azure, AWS, Redis, SQL, Postgres, PowerBI, Databricks, Snowflake, SparQL, Langchain, Langgraph, Vector Databases

---

## Publications

---

1. Bao Y, Shah AP, Narang N, Rivers J, Maksey R, Guan L, Barrere LN, Evenson S, Basole R, Miao C, Mehta A. *Harnessing Business and Media Insights with Large Language Models*. arXiv preprint arXiv:2406.06559. 2024 Jun 2.
2. Koustav Ghosal, Aljosa Smolic. *Image Aesthetics Assessment Using Graph Attention Network*, International Conference on Pattern Recognition (ICPR 2022)
3. Sebastian Lutz, Richard Blythman, Koustav Ghosal, Matthew Moynihan, Ciaran Simms, Aljosa Smolic. *Jointformer: Single-Frame Lifting Transformer with Error Prediction and Refinement for 3D Human Pose Estimation*, International Conference on Pattern Recognition (ICPR 2022)
4. Ojasvi Yadav, Koustav Ghosal, Sebastian Lutz, and Aljosa Smolic. *Frequency-domain loss function for deep exposure correction of dark images*. Signal, Image and Video Processing (2021).
5. Koustav Ghosal, Aakanksha Rana, and Aljosa Smolic. *Aesthetic Image Captioning From Weakly-Labelled Photographs.*, Workshop In Cross-Modal Learning in Real World, International Conference on Computer Vision (ICCV) 2019, Seoul
6. Xu Zheng, Tejo Chalasani, Koustav Ghosal, Sebastian Lutz, Aljosa Smolic. *STaDA: Style Transfer as Data Augmentation.*, 14th International Conference on Computer Vision Theory and Applications, VISAPP 2019, Prague
7. Koustav Ghosal, Mukta Prasad, Aljosa Smolic, *A Geometry-Sensitive Approach for Photographic Style Classification*, Irish Machine Vision and Image Processing Conference, August 2018 (IMVIP), Belfast.
8. Koustav Ghosal, Ameya Prabhu, Riddhiman Dasgupta, Anoop M. Namboodiri, *Learning Clustered Subspaces for sketch-based Image Retrieval*, Asian Conference on Pattern Recognition (ACPR), November 2015, Kuala Lumpur, Malaysia.
9. Koustav Ghosal, Anoop M. Namboodiri, *A Sketch-Based Approach to Video Retrieval using Qualitative Features*, Proceedings of the Ninth Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP), 14-17 Dec 2014, Bangalore, India.
10. Sanchit Aggarawal, Koustav Ghosal, Pulkit Singhal, Priyanka Srivastava, *Effect of Learning on Audio Spatial Working Memory*, Spatial Cognition 2014, Bremen, Germany, 15-19 September 2014.