Saurav Jha

6666 Rue Saint-Urbain, Montréal, QC H2S 3H1, Canada

www.sauravjha.com.np | saurav.jha@mila.quebec | (263) 255 7738 | srvjha

EDUCATION

University of New South Wales — **Sydney, Australia**
Ph.D. in Computer Science — *Feb 2022 - Aug 2025*
Thesis: "Enhancing continual learning of deep neural network architectures for real-world relevance".[pdf]
Advisors: Lina Yao, Dong Gong.
Teaching: COMP9418 (Statistical Machine Learning), ZZEN9444 (Deep Learning), COMP6713 (NLP).

University of St Andrews — **St Andrews, UK**
Erasmus Mundus Joint MSc in Advanced Systems Dependability — *Sep 2019 - Jul 2021*
Grade: 78.55% (First Class Honours).
Thesis: "Learning Domain-Specific Language Models for Automatic Speech Recognition". [Report][Poster]
Advisors: Imran Sheikh, Emmanuel Vincent.

Motilal Nehru National Institute of Technology Allahabad — **Prayagraj, India**
B. Tech., Computer Science and Engineering — *Jul 2014 - May 2018*
Thesis: "Cartoon Face Detection and Recognition". [Report][Code][arXiv]
Advisor: Suneeta Agarwal.

EXPERIENCE

MILA — **Montréal, Canada**
IVADO Postdoctoral Fellow, Advisor: Sarath Chandar — *Sep 2025 – Present*
- Research on test-time scaling, model compression, and world models.
- Leading industry collaboration projects with Samsung (MoE-LLMs) and Amgen (protein-LLMs).

OpenStream Inc. — **Melbourne, Australia**
Applied AI Scientist, Manager: Gholamreza Haffari — *May 2025 – Aug 2025*
- Developed Enterprise-class neuro-symbolic conversational agents for insurance underwriting in an agile environment.

LightSpeed Studios, Tencent — **Sydney, Australia**
AI Research Intern, Mentor: Shengju Qian — *Sep 2024 - Mar 2025*
- Worked on controllable image generation and preference optimization for multi-modal LLMs.

Creative AI Lab, Sony Group Corporation — **Tokyo, Japan**
Research Scientist Intern, Mentor: Shiqi Yang — *May 2024 - Aug 2024*
- Worked on continual personalization of pre-trained text-to-image diffusion models.

French Institute for Research in Computer Science & Automation (INRIA) — **Nancy, France**
Intern, European Union's COMPRISE project — *Mar 2021 - Jul 2021*
- Worked with Emmanuel Vincent on learning domain-specific language models for speech recognition.

FactSet Research Systems Inc. — **Hyderabad, India**
Machine Learning Engineer, Mentor: Keval Dave — *Jun 2018 – Jul 2019*
- Worked on incorporating deal extraction and acronym ambiguity resolution into FactSet's named entity recognition service; improved the F1-score by $\approx 6.5\%$.

SERVICES

- Co-organizer for Efficient Visual Generation Workshop, ECCV 2026. [Webpage]
- Co-organizer for the CoLLAS monthly seminar series [Webpage] and the Mila AI Colloquium.
- Reviewer for WACV 2026, ICLR 2024/5/6, CVPR 2024/5/6, NeurIPS 2023/4/6, IEEE TPAMI, TNNLS.
- Program Committee member for ECAI 2025, WWW 2025/6 (Industry track), CIKM 2023 (Workshop proposals).

TALKS AND PRESENTATIONS

- "Introspective updates for Continual learning" - Visual Computing Seminar, MIT CSAIL (March 2026).
- "Uncertainty-aware Continual Learning" - Dr. Marinka Zitnik's group, Harvard University (Virtual, April 2025).
- "On Continual Learning" - Dr. Bishesh Khanal's group, NAAMII, Patan, Nepal (May 2025).
- "Neural Processes for Continual Learning" - EEML summer school, Novi Sad, Serbia (July 2024).
- "CLAP4CLIP" - 2nd Bayes duality workshop, RIKEN AIP, Tokyo, Japan (June 2024).

HONORS AND AWARDS

- **IVADO postdoctoral fellowship 2025 (CAD 200,000)** – among the 11 researchers awarded worldwide.↗
- **Tiktok-sponsored Best Student Presentation** award at the 2024 Sydney AI meetup.
- **Best runner-up paper awardee** at the CVPR 2022 workshop on Continual Learning.
- **Best master's thesis award** for Erasmus+ DEPEND 2021 class (out of 46 students across 3 universities).
- **Best students' poster (1,000 euros prize)** at Digital Ethics4EU 2021 workshop, TU Dublin. [Poster]
- **Winner of Barclays chatbot challenge** at Hack the Burgh 2020, the University of Edinburgh. [Project]
- **2019 Erasmus Mundus scholarship (49,000 euros)** for joint Master's degree studies in the UK and France.

SCHOLARLY WORKS

Pre-prints

- **Jha S.**[*], Nilaksh[*], Zholus A.[*], Chandar S. "Reconstruction or Semantics? What Makes a Latent Space Useful for Robotic World Models", 2026. [arXiv] [Project] [Code] (* Equal Contribution)
- **Jha, S.**, Hashemzadeh, M., Pasand, A.S., Parviz, A., Lee, M., Knyazev, B. "REAM: Merging Improves Pruning of Experts in LLMs", 2026. [arXiv] [Code] [Models] [sub-reddit] **[Work with Samsung]**

Conferences

- **Jha S.**, Yang S., Ishii M., Zhao M., Simon C., Mirza J., Gong D., Yao L., Takahashi S., Mitsufuji Y., "Mining Your Own Secrets: Diffusion Classifier Scores for Continual Personalization of Text-to-Image Diffusion Models", ICLR 2025. [arXiv] [Project] **[Work with Sony]**
- **Jha S.**, Gong D., Yao L., "CLAP4CLIP: Continual LeArning with Probabilistic finetuning for Vision-Language Models", NeurIPS 2024. [arXiv] [Code]
- **Jha S.**, Gong D., Zhao H., Yao L., "NPCL: Neural Processes for Uncertainty-Aware Continual Learing". NeurIPS 2023. [Paper] [Code] [Blog]
- Li Y., Liu Z., **Jha S.**, Cripps S., Yao L., "Distilled Reverse Attention Network for Open-world Compositional Zero-Shot Learning". ICCV 2023. [Paper]

Journals

- Mirza, M. J., Zhao, M., Mao, Z., Doveh, S., Lin, W., Gavrikov, P., Dorkenwald, M., Yang, S., **Jha, S.**, Wakaki, H., Mitsufuji, Y., Possegger, H., Feris, R., Karlinsky, L., Glass, J., "GLOV: Guided Large Language Models as Implicit Optimizers for Vision-Language Models", TMLR 2025. [Project]
- **Jha S.**, Schiemer M., Zambonelli F., Ye J. "Continual learning in sensor-based human activity recognition: An empirical benchmark analysis". Information Sciences, 2021. [Paper] [Code]
- Ye J., Nakwijit P., Schiemer M., **Jha S.**, Zambonelli F. "Continual Activity Recognition with Generative Adversarial Networks". ACM Transactions on Internet of Things (TIOT), 2021. [Paper] [Code]
- **Jha S.**, Sudhakar A., Singh A.K. "Learning cross-lingual phonological and orthographic adaptations: a case study in improving neural machine translation between low-resource languages". Journal of Language Modelling, 2019. [Paper] [Code]