Court File No. T-1717-24

FEDERAL COURT

B E T W E E N:

SAMUELSON-GLUSHKO CANADIAN INTERNET POLICY AND PUBLIC INTEREST CLINIC

Applicant

- and -

ANKIT SAHNI

Respondent

AFFIDAVIT OF RAGHAV GUPTA

I, Raghav Gupta, residing in the city of San Jose, California, USA, SWEAR AND SAY THAT:

1. I am the creator of the generative artificial intelligence (AI) application called RAGHAV Artificial Intelligence Painting App ("RAGHAV"), and as such have personal knowledge of the following, except where stated to be based on information and belief in which case I do believe the same to be true.

2. I am currently employed as a Machine Learning Engineer at Cornerstone OnDemand in Dublin, California, and have extensive experience working in the field of machine learning. My current *curriculum vitae* is attached as **Exhibit A**.

How RAGHAV Works

3. RAGHAV is based on a machine learning and computer vision model, which allows it to generate stylized images in response to user input photographs or sketches.

4. One main difference between machine learning and other computer algorithms is that

machine learning algorithms learn from a vast set of rules based on the data that is fed into it. Alternatively, other computer algorithms must rely on the programmer to type in a set of predefined rules.

5. The machine learning algorithm behind RAGHAV, is based on a subfield of machine learning called neural networks (the "Network").

6. Neural Networks are programmed structures inspired from biological neurons of the nervous system.

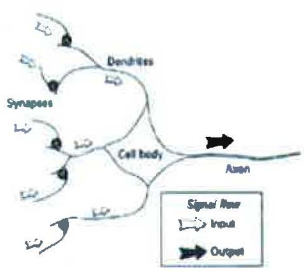

Figure 1.1 Essential components of a neuron shown in stylized form.

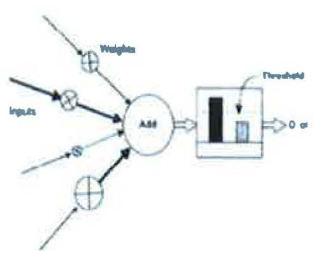

Figure 1.2 Simple artificial neuron.

Figure 1: Biological and Artificial Neurons[1]

7. A biological neuron (Figure 1, Left) takes several incoming signals through synapses, electrochemical junctions located on dendrites, branches of the neuron cell. The cell body processes all the signals and generates a resulting signal based on a threshold which gets transmitted to other neurons through the axon. Similarly, an Artificial Neuron (Figure 1, Right) takes values as inputs from multiple artificial neurons, processes them using matrix multiplications (using values called weights) and other operations, and outputs the resulting signals to other artificial neurons.

[1] Kevin Gurney, "An Introduction to Neural Networks," (London: UCL Press Limited, 1999) at 1, online: <An Introduction to Neural Networks>.

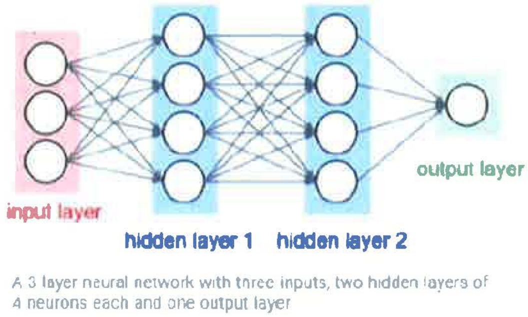

Figure 2: Neural Network[2]

8. Many artificial neurons form a layer, and many layers form a Neural Network (Figure 2). An input layer can, for example, be pixel values of an image, numerical representations of words in a text, or descriptive values in tabular data. The output layer can be a label predicting a category, such as 'dog' in an image, 'price' of a house given descriptive feature values of the house, or next word prediction given a sequence of words. The hidden layers are latent representations which form learnt intermediate features required to predict the output from given input.

9. With each pair of input and output training data provided to the neural network, the Network updates its weights in the layers. As a result, the output can be generated for the given input.

[2] "Neural Network Architectures" in *CS231n Deep Learning for Computer Vision*, online: < CS231n Deep Learning for Computer Vision>.

10. When the Network learns using all the data in training set, one hopes for the neural network to have "generalized," that is, learned enough representations to produce a correct output for any new unseen input.

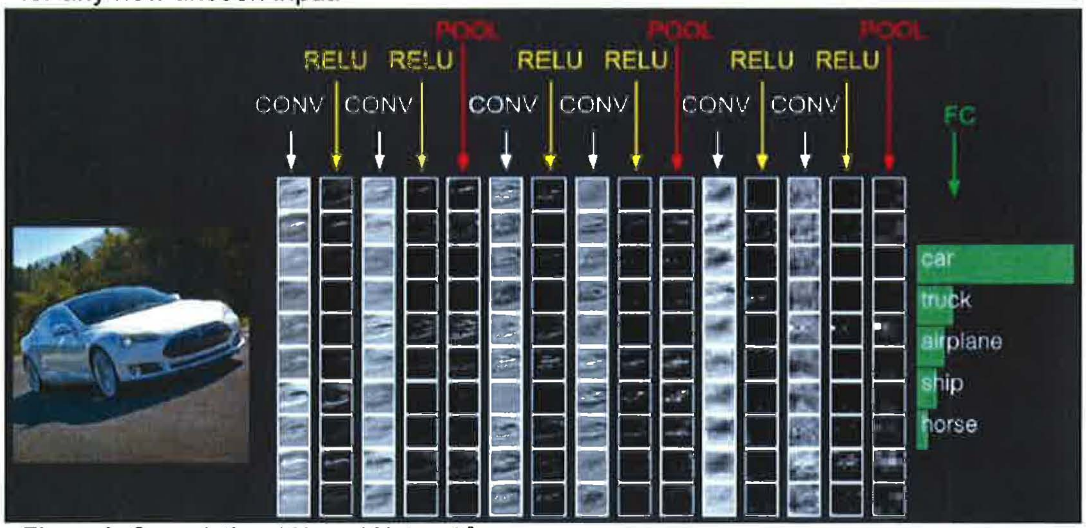

Figure 3: Convolutional Neural Network[3]

11. A Convolutional Neural Network (Figure 3) (CNN) is a neural network with efficient operations in its layers for learning features from input images. The Network learns basic features like edge detection in its initial layers, then, using these simple features, learns more complex features like textures and patterns in subsequent layers. In later layers of the network, the CNN learns features of complex objects like dog noses, human eyes and flowers. (Figure 4)

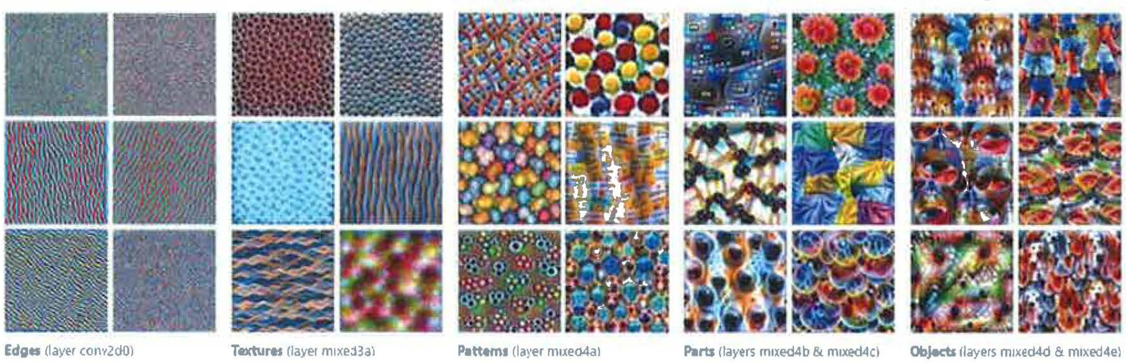

Figure 4: Features which layers learn in a CNN[4]

---

[3] "Convolutional Neural Networks (CNNs / ConvNets)" in *CS231n Deep Learning for Computer Vision*, online: < CS231n Deep Learning for Computer Vision>.

[4] Chris Olah et al., "Feature Visualization," in *Distill* (2017), online: <https://distill.pub/2017/feature-

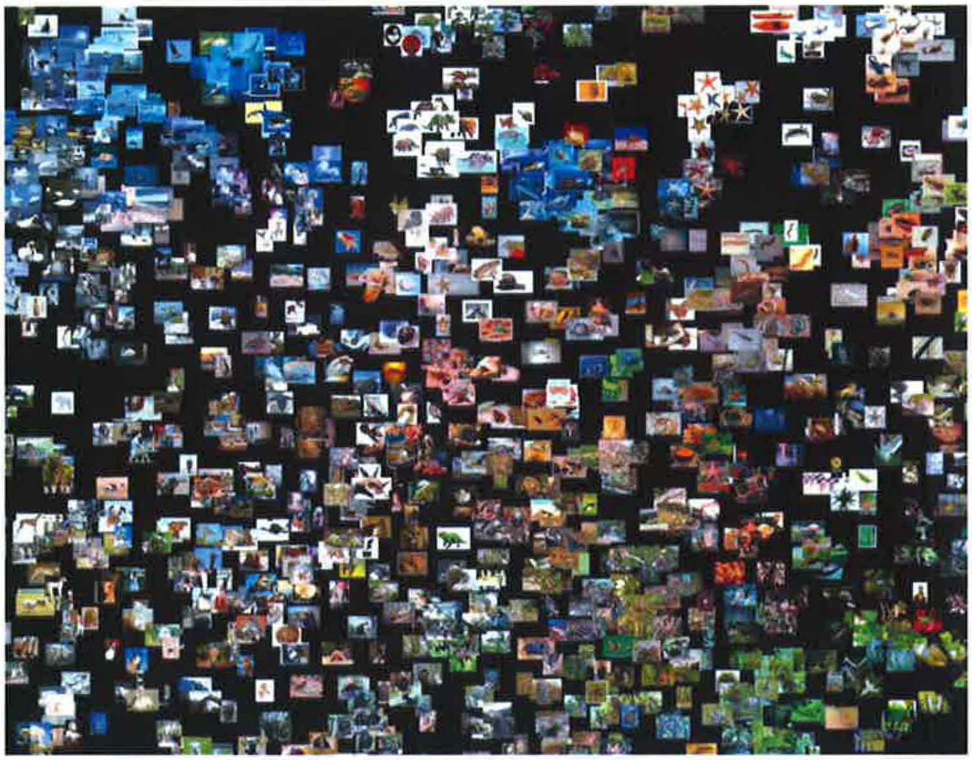

Figure 5: Feature representation space (projected to 2D space) learnt by a CNN[5]

12. A CNN can learn a multi-dimensional representation space where similar images are closer together. If we pass input images through the CNN, and extract the representational vectors from its penultimate layer, then vectors from similar images (or images belonging the same category) will cluster together. (Figure 5)

Neural Artistic Style Transfer

13. RAGHAV is based on a technique called 'neural style transfer' which is built using CNNs. Neural style transfer is a technique that allows us to generate an image with the same "content" as a base image, but with the "style" of our chosen picture.

---

visualization/>

[5] "Image t-SNE viewer," online: <Image t-SNE viewer>

14. Specifically, RAGHAV is built with a variant of neural style transfer using the research paper "Exploring the structure of a real-time, arbitrary neural artistic stylization network".[6]

15. In neural style transfer, a CNN is used to extract the features of the content and style images. (Figure 6)

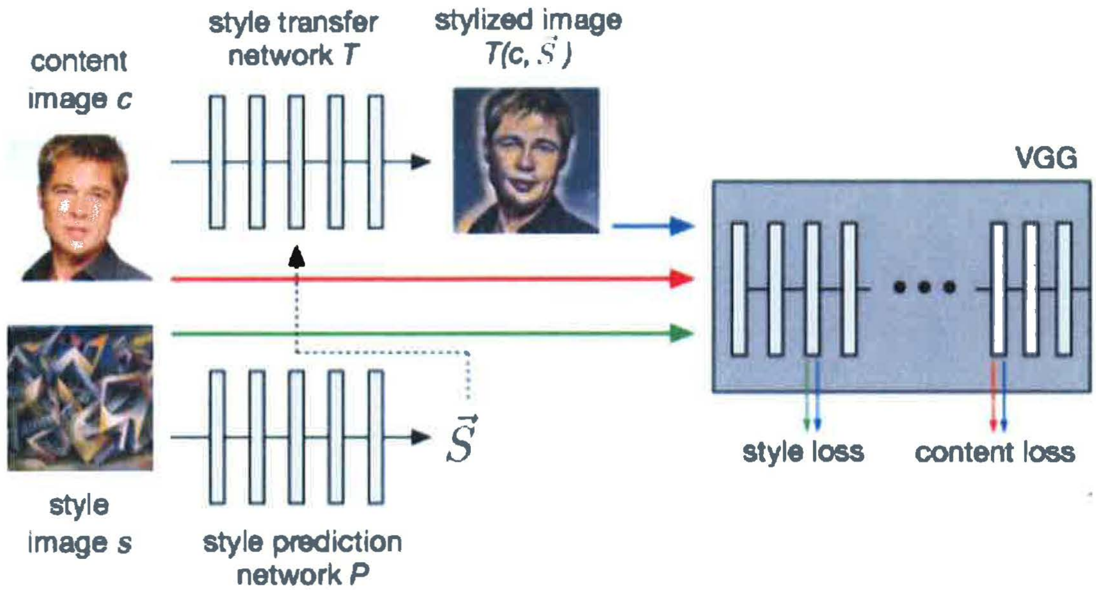

Figure 6: Architecture for Neural Style Transfer[7]

16. Neural style transfer is based on the following two propositions: (1) two images are similar in content if their high-level features as extracted by an image recognition system are close in Euclidean distance, and (2) two images are similar in style if their low-level features as extracted by an image recognition system share the same spatial statistics. This is motivated by the hypothesis that a painting style may be regarded as a visual texture. Literature suggests that repeated motifs representative of a visual texture may be characterized by lower order spatial statistics. Images with identical lower-order spatial statistics appear perceptually identical and capture a visual texture.

---

[6] Golnaz Ghiasi et al, "Exploring the structure of a real-time, arbitrary neural artistic stylization network" (Accepted as an oral presentation at British Machine Vision Conference, 2017) online: < https://arxiv.org/abs/1705.06830>
[7] *Ibid,* at 3.

17. Here, the image recognition system refers to a CNN, trained with image recognition task on a large corpus of 14M images called ImageNet. It is then trained for the task of neural style transfer with training content and style images.

18. After training, any new unseen image can be used as a style image. (Figure 7)

Figure 1: Stylizations produced by our network trained on a large corpus of paintings and textures. The left-most column shows four content images. Left: Stylizations from paintings in training set on paintings (4 left columns) and textures (4 right columns). Right: Stylizations from paintings never previously observed by our network.

Figure 7: New unseen styles can be used for Neural Style Transfer[8]

Creative Aspects of the Specific Neural Style Transfer Algorithm

19. A variable value determining the amount of style transfer between content and style can also be specified. Different values of this variable will result in different outputs. (Figure 8)

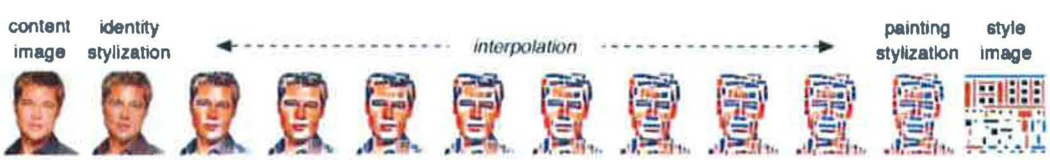

Figure 8: Linear interpolation between identity transformation and unobserved painting. Note that the identity transformation is performed by feeding in the content image as the style image.

Figure 8: Amount of Content and Style Transfer can be controlled[9]

---

[8] *Ibid,* at 2.
[9] *Ibid,* at 10.

20. Using multiple styles is also possible for a single content image by extracting and using the features from all of the style images.

21. The embedding representation space learnt by the CNN captures semantic structure of styles, that is, semantically similar styles would be clustered together.

22. The structure of the embedding representation space learnt by the CNN also permits novel exploration. The CNN model can capture a local manifold from an individual artist or painting style (Figure 9). The embedding space can be explored and new stylizations can be generated by varying local style changes for a specific painting style. Thus, new styles can be used (either entirely different or a variation of a given style image) for a different output each time for the same content image.

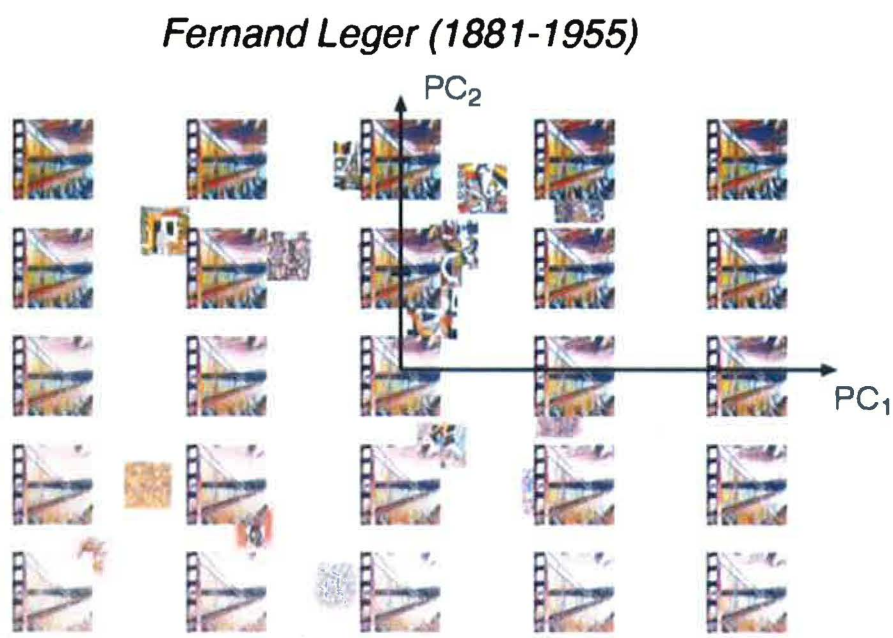

Figure 7: Exploring the artistic range of an artist using the embedding representation. Calculated two-dimensional principal components for a given artist and plotted paintings from artist in this space. The principal component space is graphically depicted by the artistic stylizations rendered on a photograph of the Golden Gate Bridge. The center rendering is the mean and each axis spans ±1 standard deviations in along each axis. Each axis tick mark indicates 2 standard deviations. Left: Paintings and principal components of Janos Mattis-Teutsch (1884-1960). Right: Paintings and principal components of Fernand Leger (1881-1955). Please zoom in on electronic version for details.

Figure 9: New styles based on an artist's artistic range can be used on the fly[10]

[10] *Ibid.*

23. As an example of the aspects of the specific neural style transfer algorithm (paragraphs 18-21), I entered the inputs used in the Suryast artwork into RAGHAV, with varying style weights and style embedding exploration for the purpose of preparing this affidavit and received the following output variations. (Figure 10).

Figure 10: Varying levels of Style Weights and Style Embedding Exploration using Suryast Inputs

<u>Response to the Affidavit of Phillip Williams</u>

24. I have reviewed the Affidavit of Phillip Williams and have the following comments.

25. In paragraph 29 of the Affidavit of Phillip Williams, Mr. Williams describes the process for generating an AI image as similar to planting a seed, stating that "while the user may choose the seed and the conditions, they do not control what grows or how it looks."

26. However, the initial conditions in which a seed is planted have a significant effect on the outcome of how the seed grows. Similarly, the initial conditions inputted in an AI model by the user have a significant impact on the image that is ultimately generated. As Mr. Williams states in paragraph 29, even "small variations in the α and β values may lead the model to produce significantly different outputs."

27. Further, the user may generate many images by varying these input parameters, and select the image produced that most closely resembles the image they were imagining.

28. Significantly though, the user may not be able to control what image the RAGHAV will generate after entering their inputs.

29. At paragraph 43 of his affidavit, Mr. Williams states that "even when a user selects particular inputs, they cannot control how the model interprets those inputs in the final output."

30. This is also true in the case of RAGHAV.

31. I have provided a detailed explanation regarding how RAGHAV functions. The user demonstrates their artistic creativity through the selection of input parameters and prediction of possible output images, as well as through selecting their desired image from a set of images that they generate. The Affidavit of Phillip Williams does not address these acts of human creativity and instead focuses primarily on the innerworkings of the AI model when generating an image. This, however, is just one aspect of the process of generating an image like SURYAST using an AI model like RAGHAV.

SWORN remotely by Raghav Gupta) stated as being located in the City of) San Jose in the State of California,) before me at the City of Toronto) in the Province of Ontario, on) May 12, 2025, in accordance with) O. Reg. 431/20, Administering) Oath or Declaration Remotely.)) ______________________) Commissioner for Taking Affidavits	Raghav Gupta ______________________ **RAGHAV GUPTA**

NOTARY PUBLIC

ONTARIO

This is Exhibit “**A**” to the Affidavit
of Raghav Gupta, sworn before
me this 12th day of May, 2025.

____

Commissioner for Taking Affidavits
LSO #: 84831V

Raghav Gupta

Linkedin://raghav-gupta-59379110b | Google Scholar:// Raghav Gupta
raghav0296@gmail.com | +1 (650) 213-7125 | San Jose, CA, US

EDUCATION

MILA (MONTREAL INSTITUTE FOR LEARNING ALGORITHMS), UNIVERSITY OF MONTREAL

Master of Science in Computer Science (Machine Learning Specialization)
Sept 2019 - July 2021 | Montreal, Canada
CGPA: 4.0/4.3

Mila is one the world's top Artificial Intelligence and Deep Learning research institutes inside University of Montreal, led by Prof. Yoshua Bengio, ACM A.M. Turing Award'18.

- First position in Data Science semester-long project competition, among 30 teams of Mila Pr.Master's students, predicting gender, age and big 5 personality traits of social media users using extracted features from image, text posts and liked pages through graphical networks (Node2Vec) and multimodal late fusion neural network.
- Other Projects: Solar irradiance (GHI) prediction using satellite imagery and efficient Tensorflow pipeline; Very low resource machine translation using iterative self-training and backtranslation.
- Elected as one of 11 Mila lab representatives.

SRM INSTITUTE OF SCIENCE AND TECHNOLOGY | B.Tech in Computer Science and Engineering

July 2014 - June 2018 | Chennai, India
CGPA: 9.52/10

- Merit Scholarship holder for all 4 years.

JOB AND RESEARCH EXPERIENCE

CORNERSTONE ONDEMAND | MACHINE LEARNING ENGINEER

Sept 2024 – Present | Dublin, CA, US
SkyHive by Cornerstone

- Managing and maintaining core ML services.
- Created novel better-performing skill to skill matching based on embedding models and cosine similarity for powering candidate-jobs matching and course recommendations.
- Implemented ML solutions for integrating SkyHive inside Cornerstone.
- Upgraded SkyHive-GPT and Ishara for answering Job Labor Market related questions and providing analysis.
- Upgraded normalization services to state-of-the-art ML techniques and model.
- Created core ML service for new Skills Passport platform - title to skills, emerging skills, skill impacted by AI, career pathways, alternate careers.

SKYHIVE | MACHINE LEARNING ENGINEER

Feb 2021 – Sept 2024 | Palo Alto, US
SkyHive got acquired by Cornerstone OnDemand.

- Managing and maintaining core ML services.
- Created and deployed a methodology for restructuring the Job Architecture of companies - successful and well-received by PwC Global and Barclays - ongoing work with other companies - large opportunity in the market.
- Created SkyHive GPT - an architecture for answering Job Labor Market related questions and providing analysis. Pretraining and finetuning Mixtral 8x7B model on our datalake - deployment through vLLM after quantization through llama.cpp. Answers can route to multiple modules - RAG, text2sql (finetuned), code generation for graphs (finetuned), graph explanation module (finetuned), Knowledge Graph based reasoning, prompt engineering using DSPy.
- Optimized our skills extraction pipeline on Databricks for large inference by various optimization methods - better hardware, optimized inference of the transformer-based models, code improvements etc. - yielded 1.5x inference speed with 2.3x reduction in costs.
- Upgraded job title normalization by creating ML Models with state-of-the-art ML techniques.
- Designed and implemented a novel methodology for automatically supporting new job categories without the need for any manual labeling. Created a custom efficient BERT pretraining pipeline with deepspeed zero optimizer and deepspeed sparse attention for long documents and very large dataset.

- Large scale deep learning model inference using PySpark.
- Remote work analysis and visualizations at micro and macro levels in the US job market for a Forrester report.
- Big data analysis using PySpark and Databricks for multiple projects.
- Upgraded job segmentation service for extracting responsibilities and qualifications from text.
- Designed and implemented a methodology for detecting uncategorized emerging titles in the job market. 18 new titles / year discovered with 2 hours of QA.
- Created a tool for detecting gender bias in job descriptions.
- Used LIME and SHAP for model explanations.
- Created a Blue Economy analysis report and contributed to an Ethical AI report.
- Created a big data analysis pipeline to create a monthly report on the current state of the job labor market - used every month till present.
- Implemented a big data pyspark pipeline on Databricks for language detection + language translation using Meta's SeamlessM4T - integrated into the main data processing pipeline.
- Compared different LLM models for normalization of job labor market entities. Fine-trained a model using PEFT to improve normalization.
- Created and deployed a new product - Career Pathways - integrating data from job transitions data + seniority parsing + title-skills.
- Hierarchical clustering of job labor market entities through Topic Modelling and BERTopic. Improved the library with creating a custom methodology to name all the hierarchical clusters.

NEXUS ROBOTICS | MACHINE LEARNING ENGINEER

May 2020 – January 2021 | Montreal, Canada

- Internship part of Master's program, received 2nd highest score (98.33%) among other 92 master's internship reports.
- Implemented state-of-the-art Object Detection and Panoptic Segmentation architectures and self-supervised learning for crop-weed segmentation and stem detection.
- Implemented Facebook's DETR, DeepLabv3+, Graph Representation Learning methods - Dual-Seg and GALD, Mask-RCNN and Knowledge Distillation.
- Internship funded and peer-reviewed by Mitacs Accelerate.

NIMBLEBOX.AI | TECHNICAL CO-FOUNDER

July 2018 – July 2019 | Chennai, India

NimbleBox is a Techstars Montreal backed startup which provides easy access to Deep Learning infrastructure, and student monitoring, performance and completion rate boost for training institutes.

- Worked as a Full-Stack developer, developing a React.js web app and contributing in Flask server with SQL DB.
- Graduated from Y Combinator Startup School'18. Mentored by Chris Field, Co-founder at Clerky, Inc.
- Helped increase growth by a monthly revenue of $10k, 1500+ users, 6 training institutes clients, & collab with 2 Learning Management System providers.
- NimbleBox awarded 1st position on pitch day at Draper University Entrepreneurship program.

BRIGHAM AND WOMEN'S HOSPITAL, HARVARD MEDICAL SCHOOL | RESEARCH INTERN

July 2017 – January 2018 | Cambridge, Massachusetts, US

Research intern at Shafiee Lab

- Core team member in creating point-of-care diagnostic devices involving optical electronic components and microfluidic components, implementing computer vision algorithms, performing statistical analysis, and analysis visualizations.
- Primarily worked on 1. Predicting embryo quality and selection during In Vitro Fertilization; 2. Predicting sperm morphology; 3. Ovulation testing using saliva images from smartphone-based optical system;
- Also created and deployed iOS & Android apps for clinicians to label images required for training DL models.
- 4 journal publications - 1 in Nature Biomedical, 1 in Fertility and Sterility, and 2 in Lab on a Chip, Royal Society of Chemistry; 8 conference papers (3 oral & 5 abstract) in Fertility and Sterility, American Society for Reproductive Medicine.

NEXT TECH LAB | RESEARCHER AND FOUNDING MEMBER

Feb 2016 – June 2018 | SRM, Chennai, India

Next Tech Lab works on ML, AR/VR, IoT, Electrical Sys. & Comp. Biology, won the Silver Award in Student Led Innovation domain at QS & Wharton School Reimagine Education'17; Current Stats: 160 students, 15 publications,

18 competition wins.
• Research in Machine Learning, CNNs, RNNs; Developed AR&VR interfaces.
• Published research paper in IEEE EEEIC/I&CPS Europe, on detection of power theft using LSTMs and smart meters.
• Re-implemented a project 'Coloring black and white world using Deep Neural Nets' to auto-colorize grayscale scenery images using CNNs and hypercolumns.
• Char-by-char generation of rap lyrics using LSTMs & CMU pronunciation dict, trained on Eminem & Dr.Dre lyrics.
• Mentored students new to ML & guided them in projects. Helped manage the lab and members.

PERSONAL AWARDS & ACHIEVEMENTS

• Created the AI software which became Canada's **first Registered AI Author in Copyright**, successfully challenging current copyright laws regarding AI authorship, garnering significant international attention and media coverage.
• Won **first position** in MIT Media Lab Reality Virtually 2017 Hackathon under Best Medical/Healthcare Hack, creating Triloka, a VR simulation for Clinical Depression with an integrated AI speech bot.
• My article on Triloka project was **featured** on BWH Next Gen and the hackathon win on BWH Awards and Honors.
• Won **first position** in Smart India 2017 Hackathon under Ministry of Steel, by building a software for detecting power theft using DL (LSTMs) and smart meters.
• Won **first position** in ITC Infotech iTech 2016 Hackathon, by creating NoMi, an AR shopping app powered with a DL (LSTM) recommendation engine.
• Joint author of **registered copyright** 'Deep Shopping Assistant' (NoMi) with the Copyright Office, Govt. Of India.

SKILLS

PROGRAMMING LANGUAGES

Python, PySpark, Javascript, C++, Java, SQL, Swift 3

TOOLS & PROGRAMMING LIBRARIES

Pytorch, Tensorflow, Keras, Sklearn, OpenCV, Flask, SQLAlchemy, React.js, PySpark, Databricks, React Native, Unity 3D, Android Studio, Xcode, Beautiful Soup 4, Git, Gimp - GNU Image Manipulation Program

LANGUAGES

English, Hindi, French (Elementary)

JOURNAL PUBLICATIONS

1. ADAPTIVE ADVERSARIAL NEURAL NETWORKS FOR THE ANALYSIS OF LOSSY AND DOMAIN-SHIFTED DATASETS OF MEDICAL IMAGES | June '21
Published in Nature Biomedical; Manoj K Kanakasabapathy, Prudhvi Thirumalaraju, Hemanth Kandula, Fenil Doshi, Anjali Devi Sivakumar, Deeksha Kartik, Raghav Gupta, Rohan Pooniwala, John A. Branda, Athe M. Tsibris, Daniel R. Kuritzkes, John C. Petrozza, Charles L. Bormann Hadi Shafiee

2. PERFORMANCE OF A DEEP LEARNING BASED NEURAL NETWORK IN THE SELECTION OF HUMAN BLASTOCYSTS FOR IMPLANTATION | Sept '20
Published in eLife; Charles L Bormann, Manoj K Kanakasabapathy, Prudhvi Thirumalaraju, Raghav Gupta, Rohan Pooniwala, Hemanth Kandula, Eduardo Hariton, Irene Souter, Irene Dimitriadis, Leslie B Ramirez, Carol L Curchoe, Jason Swain, Lynn M Boehnlein, Hadi Shafiee

3. CONSISTENCY AND OBJECTIVITY OF AUTOMATED EMBRYO ASSESSMENTS USING DEEP NEURAL NETWORKS | Apr '20
Published in Fertility and Sterility; Charles L.Bormann, Prudhvi Thirumalaraju, Manoj Kanakasabapathy, Hemanth Kandula, Irene Souter, Irene Dimitriadis, Raghav Gupta, Rohan Pooniwala, Hadi Shafiee

4. DEVELOPMENT AND EVALUATION OF INEXPENSIVE AUTOMATED DEEP LEARNING-BASED IMAGING SYSTEMS FOR EMBRYOLOGY | Nov '19
Published in Lab On A Chip; Manoj Kanakasabapathy, Prudhvi Thirumalaraju, Charles L.Bormann, Hemanth Kandula, I. Dimitriadis, I. Souter, V. Yogesh, Sandeep Pavan, Divyank Yarravarapu, Raghav Gupta, Rohan Pooniwala, Hadi Shafiee

5. AN INEXPENSIVE SMARTPHONE-BASED DEVICE FOR POINT-OF-CARE OVULATION TESTING | Nov'18
Published in Lab on a Chip, Royal Society of Chemistry; Potluri, V., Kathiresan, P., Kandula, H., Thirumalaraju, P., Kanakasabapathy, M., Pavan, S., Yarravarapu, D., Soundararajan, A., Baskar, K., Gupta, R., Gudipati, N., Shafiee, H.

6. Detection of Non-Technical Losses of Power Using Advanced Metering Infrastructure and Deep Recurrent Neural Networks | June'17

Published in Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe); Chatterjee, S., Archana, V., Suresh, K., Saha, R., Gupta, R., Doshi, F.