

Jean de Dieu Nyandwi

jeandedi@andrew.cmu.edu / nyandwi.com / @jeande_d / Google Scholar / GitHub: nyandwi

Research Interests

Multimodal AI; vision-language models; evaluation and benchmarking; post-training; multilingual NLP; data curation.

Education

Carnegie Mellon University Starting Fall 2026
Incoming Ph.D. Student, School of Computer Science, Language Technologies Institute Pittsburgh, PA

Carnegie Mellon University Aug 2022 – May 2024
M.S. in Engineering Artificial Intelligence; GPA: 3.82/4.00 Pittsburgh, PA
Selected coursework: Multimodal Machine Learning, Introduction to Deep Learning, Introduction to Machine Learning, Trustworthy AI, Applied Stochastic Processes.

University of Rwanda Sep 2017 – Jan 2022
B.S. in Electronics and Telecommunication Engineering Kigali, Rwanda

Nyanza Technical School 2014 – 2016
Advanced Level Diploma in Electronics and Telecommunication Rwanda
Ranked first nationally in Electronics and Telecommunication on Rwanda's 2016 Advanced Level National Examination.

Publications

TimeBlind: A Spatio-Temporal Compositionality Benchmark for Video LLMs
Baiqi Li, Kangyi Zhao, Ce Zhang, Chancharik Mitra, **Jean de Dieu Nyandwi**, Gedas Bertasius
arXiv, 2026 / Project page

CulturalGround: Grounding Multilingual Multimodal LLMs With Cultural Knowledge
Jean de Dieu Nyandwi, Yueqi Song, Simran Khanuja, Graham Neubig
EMNLP 2025, Main Conference, Oral / Project page

Pangea: A Fully Open Multilingual Multimodal LLM for 39 Languages
Xiang Yue*, Yueqi Song*, Akari Asai, Seungone Kim, **Jean de Dieu Nyandwi**, Simran Khanuja, Anjali Kantharuban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, Graham Neubig
ICLR 2025 (* equal contribution)

NaturalBench: Evaluating Vision-Language Models on Natural Adversarial Samples
Baiqi Li*, Zhiqiu Lin*, Wenxuan Peng*, **Jean de Dieu Nyandwi***, Daniel Jiang, Zixian Ma, Simran Khanuja, Ranjay Krishna, Graham Neubig, Deva Ramanan
NeurIPS 2024 (* equal contribution)

Research Experience

Visiting Researcher

Nov 2024 – Present

Language Technologies Institute, Carnegie Mellon University

Pittsburgh, PA

Advisor: Prof. Graham Neubig.

- Led CulturalGround, a multilingual multimodal dataset and model for grounding cultural knowledge across 42 countries and 39 languages.
- Designed a region-structured pipeline for extracting culturally grounded entities from Wikidata and generating 30M VQA samples.
- Trained CulturalPangea with region/language-aware sampling and data mixing, improving cultural benchmark performance while preserving general vision-language ability.
- Building fine-grained evaluations for chain-of-thought reasoning in VLMs, with emphasis on visual grounding, reasoning structure, and failure modes.

Graduate Student Researcher

Sep 2023 – Dec 2024

Language Technologies Institute, Carnegie Mellon University

Pittsburgh, PA

Advisor: Prof. Graham Neubig.

- Co-led NaturalBench, identifying benchmark biases that allowed “blind” solutions and building the evaluation repository for 50+ VLMs.
- Ran debiasing experiments showing that many failures came from miscalibration rather than missing visual capability.
- Contributed to Pangea, a fully open multilingual multimodal LLM for 39 languages, through evaluation across multilingual and general multimodal benchmarks.

Selected Projects

Complete Machine Learning Package / 4,900+ GitHub stars

Created an open-source ML curriculum covering Python, data analysis, classical ML, deep learning, computer vision, and NLP; widely used by learners and developers.

Understanding Compositionality in Vision-Language Models / Report

Improved VLM compositionality using synthetic hard-negative datasets and a scalable diffusion-based pipeline for vision-language data generation.

ModernConvNets / 340+ GitHub stars

Implemented and analyzed 12+ CNN architectures from AlexNet to ConvNeXt, studying design principles and performance trade-offs in visual recognition.

Teaching

- Teaching Assistant, Introduction to Machine Learning for Engineers (18-661), Carnegie Mellon University, Spring 2024.
- Teaching Assistant, Introduction to Deep Learning (11-785), Carnegie Mellon University, Fall 2023.
- Instructor, Introduction to Deep Learning, Mbaza NLP, 2022.
- Instructor, Deep Learning Bootcamp, The Python Academy, 2021.

Leadership and Service

- Conference reviewer: NeurIPS 2025 workshop; ICLR 2023 and 2024 Tiny Papers Track.
- Conference organizer: IndabaX Rwanda 2023 and 2024, Rwanda's largest machine learning gathering.
- President, Data Science Club at CMU, 2022–2023; organized ML research events, hackathons, and LLM-focused coding challenges.
- Program Associate, African Students' Education Fund, 2017–2021; led trimester meetings and school visits, co-led graduate programs, and created the organization's first newsletter series.
- Volunteer, African Students' Education Fund, 2016–2017; supported student communication, government partner communication, and student performance reviews.

Science Communication

- Creator, Deep Learning Revision, a research publishing platform for technical AI articles; published "The Transformer Blueprint" with 20K+ views.
- DeepLearning.AI Ambassador, 2020–2022; organized AI events in Kigali and taught local workshops.
- Built an X/Twitter audience of 46K+ followers by explaining AI research to a broad technical audience.

Professional Experience

Data Scientist

Jan 2021 – Feb 2022

VIEBEG Technologies Ltd

Kigali, Rwanda

- Analyzed sales data to improve inventory management and built real-time analytical dashboards; contributed to a 20%+ revenue increase.

Selected Honors

Ranked first nationally in Electronics and Telecommunication on Rwanda's 2016 Advanced Level National Examination; DeepLearning.AI Ambassador Spotlight, 2022; Top 50 AI Influencers, Onalytica, 2022; African Students' Education Fund Scholar, 2014–2016.

Technical Skills

Languages: Python, MATLAB, Bash, LaTeX.

Frameworks: PyTorch, TensorFlow, Hugging Face, Detectron2.

Tools: Git, Weights & Biases, NumPy, Pandas, Scikit-learn, Matplotlib.