

5th Workshop on practical ML for limited/low resource settings (PML4LRS) @ ICLR 2024

Workshop Summary

The constant and breakneck speed of progress being made in artificial intelligence (AI) and generative AI needs to be resource optimized for practical societal impacts. Adapting the state-of-the-art (SOTA) methods such as large language models (LLMs), Diffusion Models, and Neural Radiance Fields (NeRFs) to resource-constrained environments, to run (even few-shot fine-tuning and inference) under low resources such as those typically in developing countries and computing at the edge, is highly challenging in practice. Partly due to the lack of diversity in the data and personnels (involved in annotating and validating), their high demand in computational resources, variations in the selection of performance metrics. For example, recent breakthroughs in natural language processing (NLP), computer vision, and speech analysis, for instance, rely on increasingly complex and large models (e.g. most models based on transformers and attention such as BERT, GPT-2/GPT-3, DALLE-2, and stable diffusion) that are pre-trained in on large corpus of unlabeled data. Applying these models in a resource-constrained environment is a non-trivial challenge. Moreover, the potential risks associated with such large models in low resource settings, e.g., disinformation, is virtually unexplored. Low/limited resources mean a hard path towards the adoption of these breakthroughs for most edge applications as well as in developing countries. As a result, most of these advances are limited to giant technology companies and institutions that have access to computational resources and big datasets which, unintentionally, marginalizes institutions and companies with lower resources as well as significantly hampers edge use cases. These challenges downgrade the overall trustworthiness of such AI solutions to achieve positive societal impact worldwide. Methods such as data-augmentation, transfer learning and synthetic data will not solve the problem either due to bias in the original pre-training datasets as well as the prohibitive cost and resource needs of fine-tuning these large scale models. If disparities in resources continue as models become more resource-intensive, it would exacerbate the widening income inequalities across the globe and disturb the constant progress towards equity.

Recent progress with focus given to ML for social good has the potential to alleviate the problem in part. However, the themes in such workshops are usually application-driven such as ML for healthcare and ML for education, and less attention is given to practical aspects as it relates to developing countries in implementing these solutions in low or limited resource scenarios. This, in turn, hinders the democratization of AI in developing countries.

As a result, we aim to fill the gap by bringing together researchers, experts, policymakers, and related stakeholders under the umbrella of practical ML for developing countries that aim to facilitate design and adoption of trustworthy ML solutions with embedded flexibility for low-resource settings. The workshop is geared towards fostering collaborations and soliciting submissions under the broader theme of practical aspects of implementing machine learning (ML) solutions for problems in developing countries. We specifically encourage contributions that highlight challenges of learning under limited or low resource environments that are typical in developing countries. Among the specific themes include, but not limited to: learning from limited labeled data, selective augmentation and evaluation of synthetic data from generative models, light model design, training and validation, data-centric AI techniques, active learning,

data labeling with a limited budget (such as model-based pre-labeling), parameter/filter pruning (such as Gate Decorator), model compression and quantization, efficient algorithms under extreme resource constraints (such as active learning and pre-labeling with pre-trained models), learning from weak labels, and policy issues of applying ML in developing countries. We believe that such methods would be able to enable continued progress in ML while also being inclusive of those with limited resources. It also helps lower the carbon footprints of training these increasingly complex models. Recent reports, for instance, indicated that the carbon footprint of training GPT-3 is equivalent to lifetime emissions by an automobile.

We conducted the [1st workshop on Practical ML for Developing Countries at ICLR 2020](#), followed by subsequent iterations in ICLR 2021 (2nd workshop), [ICLR 2022](#) (3rd workshop), and ICLR 2023 (4th workshop). Throughout these iterations the workshop attracted a growing number of participants, mainly from the African continent, and hosted a number of top-researchers, practitioners and policy makers from academia, industries and non-governmental agencies. By expanding our thematic areas relevant to developing countries and algorithms with less footprints, we intend to attract even more audience in this fourth installment.

Topics: the scope of the workshop includes the following broad areas:

Algorithms and Methods

- Methods for collecting and generating training data within *data-scarce* (limited labeled data) settings (such as weak labels, model-based pre-labeling, teacher-student models, and transfer learning).
- Machine learning techniques applied to limited data (e.g. active learning, few-shot, and zero-shot learning).
- Approaches to training and inference on resource-constrained devices with very low footprints (such as model quantization, model compression, model distillation, low precision training, model pruning methods, and generalized model optimizations).
- Alternative learning methods coupled with deep models targeted for low resource settings.
- Self-supervised and semi-supervised approaches in the absence of high quality training data.
- Analyze models in the perspective of fairness, explainability and mitigation strategy in both dataset creation as well as algorithmic solutions.
- Automated techniques to stratify and evaluate data quality and fair representation in order to increase throughput in low-resource settings.
- Ensure generalizability of ML in limited resource settings by injecting domain knowledge.
- Effectively integrate multiple data sources, which are often in silos.
- Synthetic training data.

Industry Experience and Applications

- Data science and engineering practices that help balance accuracy/latency tradeoffs while scaling ML models in low resource environments.
- Model deployment and monitoring temporal drift.

-
- Measuring success or impact that goes beyond algorithmic metrics (such as accuracy or F1 score).
 - Data-centric techniques that support public institutions (government transparency, healthcare, education, etc).
 - High value applications relevant to developing countries such as climate change, natural disasters, conflict resolution, and tackling malnutrition.
 - Large language models and their associated risks in low resource settings.

Social and Policy Topics

- Successful AI solution implementation stories which work at a small scale (e.g. local institution, city) that could be applied at a larger scale.
- Connecting skilled professionals with organizations that deeply understand the local problems.
- Securing funding for proof-of-concept (POC) projects or for scaling existing POCs.
- Building effective research and implementation teams, with a focus on challenges specific to developing regions such as countries in Africa.
- When machine learning is NOT a viable option.
- Strategies and policies enabling or enhancing AI/ML adoptions for developing countries.

Submission types

- Short papers and position pieces (up to 5 pages)
- Problem statements and abstracts (up to 2 pages)

5-page submissions will be *eligible* for oral or poster presentations. 2-page submissions will be presented as posters.

Tentative Schedule

The workshop will be a full-day event from 9:00 am - 6:15 pm on May 11th, 2024 in Vienna, Austria. The tentative schedule is as follows:

- 09:00 - 09:10 Welcome Remarks
- 09:10 - 09:45 Invited Talk 1
- 09:45 - 10:00 Contributed Talk 1
- 10:00 - 10:15 Contributed Talk 2
- 10:15 - 10:30 Contributed Talk 3
- 10:30 - 11:00 Coffee Break + Poster Session I
- 11:00 - 11:35 Invited Talk 2
- 11:35 - 11:50 Contributed Talk 4
- 11:50 - 12:05 Contributed Talk 5
- 12:05 - 12:20 Contributed Talk 6
- 12:20 - 02:15 Lunch + Poster Session II
- 02:15 - 02:30 Contributed Talk 7
- 02:30 - 02:45 Contributed Talk 8
- 02:45 - 03:00 Contributed Talk 9
- 03:00 - 03:15 Contributed Talk 10
- 03:15 - 03:45 Coffee Break + Poster Session III
- 03:45 - 04:20 Invited talk 3

-
- 04:20 - 04:55 Invited Talk 4
 - 04:55 - 05:55 Panel Discussion: Redesigning AI solutions centered on contexts in limited resource settings
 - 05:55 - 06:15 Awards and Closing Remarks

Each oral presentation will be 10 minutes followed by 5 minutes for Q/A, each invited speaker presentation will be 25 minutes followed by 10 minutes for Q/A for increased discussion. Since the idea is to bring a diverse group of participants to discuss practical issues of implementing ML for developing countries, we specifically added a panel discussion focused on the topic with panelists from representatives of policymakers, researchers, industry experts, and other stakeholders.

Invited Speakers (Candidates)

- [Kathleen Simunyu](#):

Kathleen is a data scientist who enjoys building and maintaining data infrastructure as well as discovering patterns that uncover insights from data. She works for Africa's Talking. Passionate about the democratization of machine learning, she co-founded and manages a data science and machine learning community which focuses on encouraging individuals, with a special focus on women, to get into the field. With over 1000 members, the community is designed to help individuals interested in growing their skills by supporting their learning journeys, connecting them with peers for collaboration as well as connecting them to opportunities for work. Through the connections fostered and the work the community is doing, we have partnered with Google through the Together with Google Developers program in Sub Saharan Africa, Africa's Talking Ltd, Moringa School, Intel, BRaVe Ventures and NVIDIA AI among others. Kathleen has been recently selected as one of MIT's 35 under 35 Innovators 2022.

- [Dalton Lunga](#):

Dalton is currently a senior research scientist in machine learning-driven geospatial image analytics and a group leader for GeoAI at ORNL. In this role, he deploys machine learning and computer vision techniques in high-performance computing environments, focusing on creating imagery-based data layers of interest to various societal problems, e.g., enabling accurate population distribution estimates, global high-resolution gravity map generation, and damage mapping for national security needs. He currently researches and develops machine learning techniques and advanced workflows for handling large volumes of geospatial data. A graduate of Purdue University and former employee at the council for scientific and industrial research in South Africa, Dalton brings both academic and industrial background to ORNL to help lead scientific discovery in geospatial challenges for societal impact.

- [Sara Hooker](#)

Sara leads Cohere For AI (C4AI) a non-profit research lab that seeks to solve complex machine learning problems, which supports fundamental research that explores the unknown, and focuses on creating more points of entry into machine learning research. Prior to For AI, Sara was a research scientist at Google Brain doing work on training models that go beyond test-set accuracy to fulfill multiple desired criteria -- interpretable, compact, fair and robust. Sara enjoys

working on research problems where progress translates to reliable and accessible machine learning in the real-world.

- [Dina Machuve](#)

Dina Machuve is the Co-Founder and CTO of DevData Analytics. She was until recently a Senior Lecturer and Researcher at the Nelson Mandela African Institution of Science and Technology (NM-AIST) in Arusha, Tanzania. Her research focuses on developing data-driven solutions in Agriculture. For her PhD, she investigated the information logistics of small and medium size food processors. Her most recent project focused at developing a poultry diseases diagnostics tool using computer vision and bioinformatics methods for small and medium scale poultry farmers in Tanzania. She was an Early Career Fellow of the Organization for Women in Science for the Developing World (OWSD). She serves in the board of Data Science Africa (DSA), a non-profit organization that builds capacity of data science and machine learning in Africa. She is also a member of a research group on Mechanism Design for Social Good (MD4SG). She completed her PhD in Information and Communication Science and Engineering from NM-AIST in 2016. She also has an MS in Electrical Engineering and BSc in Electrical Engineering.

- [Joshua Blumenstock](#)

Joshua Blumenstock is a Chancellor's Associate Professor at the U.C. Berkeley School of Information and the Goldman School of Public Policy. He is the Co-director of the Global Policy Lab and the Center for Effective Global Action. Blumenstock does research at the intersection of machine learning and empirical economics, and focuses on using novel data and methods to understand the causes and consequences of global poverty, and to improve the lives of disadvantaged people around the world. He has a Ph.D. in Information Science and a M.A. in Economics from U.C. Berkeley, and Bachelor's degrees in Computer Science and Physics from Wesleyan University. He is a recipient of awards including the NSF CAREER award, the Intel Faculty Early Career Honor, and the U.C. Berkeley Chancellor's Award for Public Service. His work has appeared in general interest journals including Science, Nature, and Proceedings of the National Academy of Sciences, as well as top economics journals (e.g., the American Economic Review) and computer science conferences (e.g., ICML, KDD, AAAI, WWW, CHI).

Diversity commitment

We expect the workshop topic and its focus areas would attract a wide range of participants such as ML researchers, industry professionals, government stakeholders, policymakers, healthcare workers, social scientists, and educators. We expect the workshop will attract a large pool of developing countries' talent that are directly affected by the problem this workshop is aimed at solving. Moreover, this demographic is the most underrepresented in the field of AI/ML. With most of the organizers of the workshop having previous experience in organizing diversity and inclusion workshops such as Black in AI and the Deep Learning Indaba, the utmost effort will be made to attract a diverse group. Specifically, the invited speakers, oral speakers, panelists will be selected taking demographics balance into account. The organizers and program committee members have already diverse sets of backgrounds. Among the five invited speakers listed above, we made an effort to balance by gender, geography, and background. We

also already have a pool of program committee members selected for the workshop. In particular, the majority of the organizing committee members and invited speakers are females.

Accessibility

To make the workshop accessible to those that are unable to attend in person, we plan to have a remote presence for authors of accepted posters to present remotely and plan to live stream and archive recordings of the whole workshop. The organizing team has extensive experience in remote presence for multiple years at the Black in AI workshop, Deep Learning Indaba, Data Science Africa, and related events when participants were unable to attend due to visa denials, for instance.

Previous related workshops

ICLR: Practical Machine Learning for Developing Countries (PML4DC) 2020, 2021, 2022, and 2023, Learning from Limited Labeled Data (LLD); Workshop: Representation Learning for Weak Supervision; and Beyond; AI for Social Good

NeurIPS: Machine Learning for Developing Worlds (ML4D); AI for social good; Black in AI

Hosted in Africa: Data Science in Africa; Deep Learning Indaba

Plan for funds

Our organizing team has extensive experience running several workshops with extensive experience raising sponsorship funds for travel grants and other expenses. We expect to raise funds from corporate sponsors for use for giving some of our presenting participant's travel grants and other logistical expenses.

Summary of reflections

1. **Diversity:** In particular, the majority of the organizing committee members and invited speakers are females.
2. **Discussion slots:** Ample time is allocated to generate discussion in the tentative schedule of the workshop, including 15 mins for Q/A for each of the four invited speakers, combined ~3hrs for coffee break and poster sessions. We also allocated a dedicated panel discussion (~1hr) that aims to bring leading researchers, practitioners and policy makers with active Q/A with the attendees.
3. **Specific problem:** resource aware machine learning. The effectiveness of AI solutions for limited resource settings needs to be scrutinized. Particularly, as AI models are going to be more data intensive and sophisticated with limited replicability in developing countries and/or low resource settings. The organizers argue that it is a timely and fitting topic to ICLR 2023, which is happening in one of the developing countries in the world.
4. Already published works are not welcome for submission. Organizers are aiming to make the PML4DC workshop a venue where experts and junior researchers meet and generate discussions for future potential collaborations. We discourage both submission and presentation of already published machine learning works.
5. **Varying lengths and scopes.** Organizers welcome two submission types: Short papers and position pieces (up to 5 pages) and Problem statements and abstracts (up to 2 page)

-
6. **Accessibility:** To make the workshop accessible to those that are unable to attend in person, we plan to have a remote presence for authors of accepted posters to present remotely and plan to live stream and archive recordings of the whole workshop. The organizing team has extensive experience in remote presence for multiple years at the Black in AI workshop, Deep Learning Indaba, Data Science Africa, and related events when participants were unable to attend due to various reasons.
 7. **Workshop organizing experience:** The organizing team has extensive experience running several workshops including Deep Learning Indaba, Data Science Africa in the African continent and in other prior venues such as ICLR, ICML, ICCV, NeurIPS, CVPR and AAAI.

Organizers biographies

Matimba Shingange

Matimba Shingange is currently pursuing NLP-focused MIT in Big Data Science at the University of Pretoria, under the supervision of Prof. Vukosi Marivate. Matimba is an experienced data analyst with a demonstrated history of working in the telecommunications industry. Skilled in Python, Analytical Skills, Microsoft Word, C++, PL\SQL Development and Java. Strong information technology professional with a Bachelor of Science (BSc) with Honours focused in Computer science from University of the Witwatersrand, Johannesburg.

Aisha Alaagib: Aisha is passionate about Machine learning for people with disability, representation learning and multimodal learning. She is also broadly interested in Uncertainty and Robustness in Deep Learning. Aisha is a recent graduate of African Institute for Mathematical Sciences (AIMS), had successful internship placements at top machine learning research institutes including Mila - Quebec Artificial Intelligence Institute .

Waheeda Saib is a Quantum Algorithms Researcher at Leiden University and IBM Q ambassador at IBM Research-Africa in Johannesburg, South Africa. As an IT professional with over 10 years of experience, she has a proven track record in development of high-quality software systems and technical innovations. At IBM Research Africa, Data driven healthcare team, she researches and develops machine learning applications in collaboration with Research scientists and institutes to meet Africa's grand challenges in Cancer and Tuberculosis while pursuing a masters in Data Science and Informatics at the University of Edinburgh. Waheeda leads a team of researchers on a project that aims to address South Africa's and Africa's cancer reporting delay using machine learning. She was also co-organizer of this workshop since its inception.

Meareg Hailemariam is a lecturer with Msc. in computer science from Addis Ababa University. Meareg is interested in AI and specifically in its applications and potentials to efficiently address societal and ecosystem challenges and looks forward towards contributing to its universality. He was also co-organizer of this workshop since its inception.

Timnit Gebru is the founder and executive director of the Distributed AI Research Institute (DAIR), honored in the 2019 [Bloomberg 50](#) list, and was selected among the [Top 30 Women in AI by Forbes](#). Prior to that she was the co-lead of the Ethical AI team at Google until her abrupt firing for bringing up issues of workplace discrimination. She did a postdoc at Microsoft Research, New York City in the FATE (Fairness Transparency Accountability and Ethics in AI) group, where she studied algorithmic bias and

the ethical implications underlying any data mining project, and received her PhD from the Stanford Artificial Intelligence Laboratory, studying computer vision under Fei-Fei Li. Her work to boost diversity and inclusion has been substantial in conferences such as NeurIPS, ICLR, CVPR, and ICCV, to name a few. Timinit has been involved in the organization of multiple workshops and conferences and she was co-organizer of this workshop since its inception.

Esube Bekele is a Senior Computer Vision Architect at IQT. Before joining IQT, he was a [National Research Council \(NRC\)](#) fellow of the National Academies at the US Naval Research Lab working on machine learning and computer vision for surveillance applications, face recognition, person re-identification, object recognition, and cognitive systems coupled with deep learning methods such as human-level explanations and hierarchical ontologies. He received the PhD from Vanderbilt University with a research focus on robotic and virtual reality-based Autism and Schizophrenia therapy systems. He is a co-organizer of the [Black in AI](#) workshop at NeurIPS in 2018 and 2019 as sponsorship, travel grant, and logistics chair. He has served as PC/reviewer/AC for several journals, conferences and workshops such as NeurIPS, ICLR, CVPR, AAAI, AISTATS, ICRA, TBME, TNSRE, and organizer of workshops at ICCV, ICLR, and NeurIPS. He was logistics chair for ICLR 2020 and he is virtual experience chair of ICCV 2021. He also served as co-organizer of this workshop since inception. He's currently serving as associate editor-in-chief of the Pattern Recognition journal.

Kevin Compher is a Vice President of Technology at IQT Labs and Foresight Institute Fellow focused on AI Assurance, Policy and Interpretability. He's a seasoned manager, engineer and data scientist with over twenty-five years of strategic experience developing high performance analytic platforms and teams in government, industry and academia. His focused specializations include applied research, development and scaling of streaming data applications, high performance computing, FinTech/RegTech, Market Data Infrastructure, Remote Sensing Phenomenology, Modeling & Simulation, Digital Twin Applications, Deep Reinforcement & Semi-supervised Learning, as well as, Explainability, Trust and Governance in AI/ML systems. He is an annual member of the review committee for the ACM's International Conference for AI in Finance, since its inception in 2019. He was also part of the organizing committees for NeurIPS 2020 Workshop on Fair-AI in Finance, & NeurIPS 2019 Workshop on Robust AI in Financial Services.

John Wamburu is a Research Scientist at IBM Research Africa working in the AI team. His work focuses mainly on anomalous pattern detection with applications in Trustworthy AI and Subset Discovery. He holds a PhD in Computer Science from the University of Massachusetts Amherst.

Advising Committee:

Girmaw Abebe Tadesse

Girmaw is a Staff Research Scientist at IBM Research Africa working in the AI Science team, mostly on detecting and characterizing systematic deviations in data and model outputs. He is also leading cross-lab projects on representation analysis, academic materials (in collaboration with Stanford University) and performance evaluation of generative models. At IBM, Girmaw also collaborates with Bill & Melinda Gates Foundation and Harvard University on data-driven insight extraction for maternal, newborn and child health. Prior to joining IBM, Girmaw worked as a Postdoctoral Researcher at the University of Oxford, where he primarily developed deep learning techniques to assist disease diagnosis. Girmaw completed his PhD at Queen Mary University of London, under the Erasmus Mundus Double Doctorate Program in Interactive and Cognitive Environments. His PhD research focused on computer vision and machine learning algorithms for human activity recognition using wearable sensors. He has worked in various research groups across Europe, including the BarceonaTech in Spain, KU Leuven in Belgium, and INESC-ID in Portugal. Girmaw is an Executive Committee Member for IEEE Kenya Section, and currently serving as a reviewer for multiple IEEE Journals and as a program committee and/or organizer for top-tier AI conferences and workshops.

Vukosi Marivate is Chair of Data Science University of Pretoria and Deep Learning Indaba co-founder. He is based at the University of Pretoria as the ABSA Chair of Data Science. Vukosi works on developing Machine Learning/Artificial Intelligence methods to extract insights from data. A large part of his work over the last few years has been in the intersection of Machine Learning and Natural Language Processing (due to the abundance of text data and the need to extract insights). Vukosi is interested in Data Science for Social Impact, using local challenges as a springboard for research. In this area, Vukosi has worked on projects in science, energy, public safety and utilities. He was also co-organizer of this workshop since its inception.

Judy Wawira is a multidisciplinary researcher, trained as both an informatician and a clinically active radiologist. She is an assistant professor at Emory university, and works in Interventional Radiology and Informatics. She has been funded through the Grand Challenges Canada, NBIB and NSF ECCS. Her career focus is on validating machine learning models for health in real clinical settings, exploring explainability, fairness, and a specific focus on how algorithms fail. She has worked on the curation of datasets for the SIIM (Society for Imaging Informatics in Medicine) hackathon and ML committee. She volunteers on the ACR and RSNA machine learning committees to support the AI ecosystem to advance development and use of AI in medicine. She is currently working on the sociotechnical context for AI explainability for radiology, especially the dimensions of human factors that govern user perceptions and preferences of XAI systems. She has been involved in organizing the Black in AI workshop and reviewed papers for the Machine Learning for Health and Machine Learning for Developing Countries at the 2019 NeurIPS workshop. She volunteers for the WHO-ITU AI working group on machine learning for global health. She has served as co-organizer of this workshop since the 2nd workshop in 2021.