DMLR: Data-centric Machine Learning Research - Past, Present and Future

Luis Oala[1]*, Manil Maskey[2], Lilith Bat-Leah[3], Alicia Parrish[4], Nezihe Merve Gürel[5], Tzu-Sheng Kuo[6], Yang Liu[7,8], Rotem Dror[9], Danilo Brajovic[10], Xiaozhe Yao[34], Max Bartolo[11], William Gaviria Rojas[12], Ryan Hileman[13], Rainier Aliment[4], Michael W. Mahoney[14,15,16], Meg Risdal[17], Matthew Lease[18], Wojciech Samek[19,20], Debo Dutta[21], Curtis Northcutt[22], Cody Coleman[12], Braden Hancock[23], Bernard Koch[24], Girmaw Abebe Tadesse[25], Bojan Karlaš[26], Ahmed Alaa[16], Adji Bousso Dieng[27], Natasha Noy[4], Vijay Janapa Reddi[26], James Zou[28], Praveen Paritosh[29], Mihaela van der Schaar[30], Kurt Bollacker[29], Lora Aroyo[4], Ce Zhang[31,24], Joaquin Vanschoren[32], Isabelle Guyon[4,33], Peter Mattson[4,29]

[1]Dotphoton, [2]NASA, [3]Mod Op, [4]Google, [5]TU Delft, [6]Carnegie Mellon University, [7]UC Santa Cruz, [8]ByteDance Research, [9]University of Haifa, [10]Fraunhofer IPA, [11]Cohere, [12]CoactiveAI, [13]Talon, [14]UC Berkeley, [15]ICSI, [16]LBNL, [17]Kaggle, [18]UT Austin, [19]TU Berlin, [20]Fraunhofer HHI, [21]Nutanix, [22]Cleanlab, [23]Snorkel AI, [24]University of Chicago, [25]Microsoft AI for Good Lab, [26]Harvard University, [27]Princeton University, [28]Stanford University, [29]MLCommons, [30]University of Cambridge, [31]Together, [32]Eindhoven University of Technology, [33]University of Paris-Saclay, [34]ETH Zurich

Abstract

Drawing from discussions at the inaugural DMLR workshop at ICML 2023 and meetings prior, in this report we outline the relevance of community engagement and infrastructure development for the creation of next-generation public datasets that will advance machine learning science. We chart a path forward as a collective effort to sustain the creation and maintenance of these datasets and methods towards positive scientific, societal and business impact.

Keywords: data-centric machine learning, artificial intelligence, datasets, impact

1 Data Ambivalence in Machine Learning

Why state the obvious? Do we really need to emphasize some machine learning (ML) research as *data-centric*? Hasn't ML science, at its core, always been just that? After all, designing algorithms that extract models from data is machine learning's *summum bonum*. In the pursuit of this goal we often oscillate between two dominant phases: (i) design algorithm and throw data at it, (ii) go back to data (and its intermediate representations) to design better algorithm. This feedback loop informs the ambivalence towards data that many of us will encounter in machine learning practice: on the one hand, we want the algorithm to extract a model from data automatically; on the other hand, we often need to analyze the data and model manually to build good algorithms. Through the lens of this oscillation, data-centric machine learning research (DMLR) can broadly be described

∗. To get involved in the community please join the discord at `https://discord.gg/FswYXMv4j9`. For updates to the manuscript you can contact `luis.oala@dotphoton.com`.

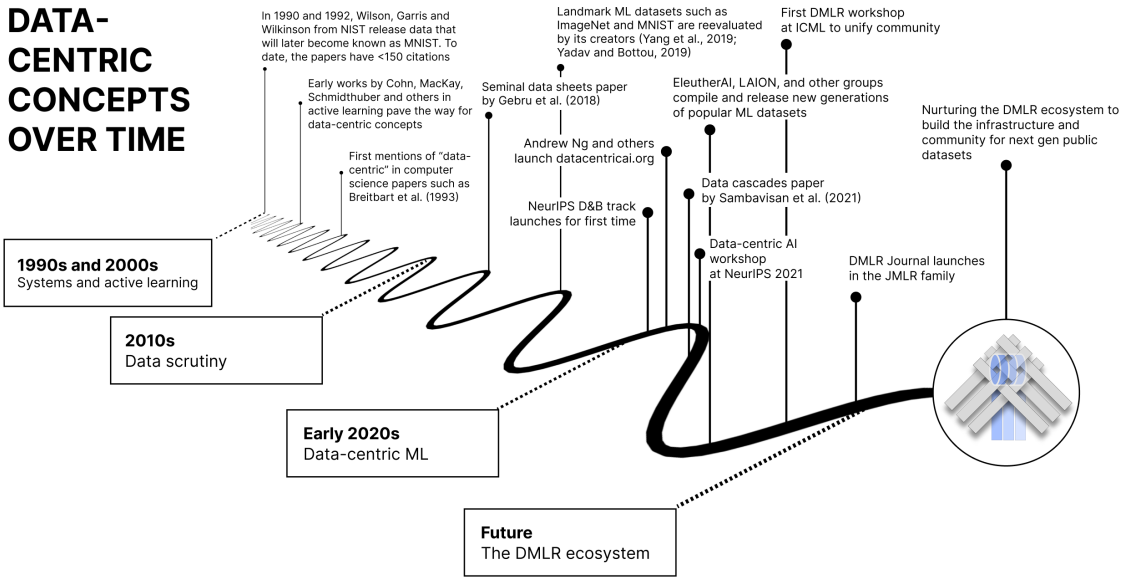

Figure 1: A timeline of some inflection points in the development of data-centric ideas.

as infrastructure, methods and communities revolving around phase (ii). In this editorial we outline key coordinates and objectives of DMLR, contextualize its origins, and summarize activities that took place at the inaugural DMLR workshop. And these lines are an invitation, a call on you, the reader, to join us in shaping the DMLR future. Be it as open source contributor, community organizer, researcher or reviewer, your ideas and efforts are needed to maintain and shape the DMLR ecosystem further.

2 Past: Data-Centricity Over Time

Historically, the ambivalence towards data has manifested in different ways. In the early 1990s, Wilson, Garris and Wilkinson [1–4] distributed "*Handwriting Sampling Forms*" at the National Institute of Standards and Technology (NIST), digitizing the resulting data into the raw ingredients that were later turned into the now infamous machine learning staple MNIST. But as of October 4, 2023, their original publications have less than 150 citations combined. In comparison, the seminal LeNet paper by [5], which is often used as stand-in reference for the MNIST dataset, sits at 60,000 citations today.[1].

This is not to open artificial fault lines à la "*data people*" versus "*model people*". But one can wonder what such artifacts reveal about the incentives in machine learning and how conducive they are to machine learning progress. Or whether, as [7] suggest, it slows progress because "*everyone wants to do the model work, not the data work*". Jumping to today, there are active and encouraging efforts in the community to counter this imbalance, most prominently the *Datasets and Benchmark Track* at Neural Information Processing System (NeurIPS) that was conceived for the first time in 2021 by Joaquin Vanschoren,

1. As an aside, LeCun et al. [5] themselves did not cite the NIST prior works [1–4]. Notably, Yadav and Bottou [6] later revisited the history of MNIST.

Serena Yeung, Maria Xenochristou. We have to emphasize that not all data work goes under-appreciated. One must only look at ImageNet [8] or CIFAR [9] for great success stories.

Around the same time as MNIST, the concept of data-centricity started to appear literally in early works by [10] and others. It was likely discussed in the systems and database circles long before the idea became an increasingly growing focus of research in the core machine learning community. The connection to systems persists to this day, evident in venues such as MLSys[2] or the DEEM workshops[3], due the high importance of optimized infrastructure to orchestrate and execute data transformation and machine learning workloads. Success stories can be found in frameworks to build and store models such as Torch [11], Theano [12], Caffe [13], TensorFlow [14], JAX [15] or PyTorch [16]. Sometimes these frameworks also led to optimized data formats, such as TensorFlow's TFRecord. Additionally, platforms like Kaggle, HuggingFace or OpenML have emerged as de-facto community data hubs and standardized data loading infrastructure. A new wave of emerging open-source projects such as Lance[4] aim to address existing gaps with respect to data loading needs. However, despite these advances, challenges regarding the compatibility, mutability, and collaborability of datasets persist. Encouragingly, new initiatives, such as 'Croissant'[5], take stabs at the Babylonian tower of data formats, uniting key stakeholders in an effort to streamline data-centric machine learning infrastructure. Similarly, DataPerf [17] is a recent community-led benchmark suite for evaluating ML datasets and data-centric algorithms, enabling the ML community to iterate on datasets, instead of just architectures.

Alongside developments in infrastructure, over time we have also witnessed critical advances in the way datasets are collected, curated and maintained. From the beginnings of modern statistical science [18; 19], active learning, a set of methods concerned with data curation, has planted its roots firmly in the machine learning and statistical learning literature [20–26]. The next generation of machine learning datasets are characterized by dense metadata annotation [27–29], collaborative refinement [30–33], user preference and human feedback [34], and evolution over time [35–37], similar to the way we treat code for computer programs already today. Data provenance and ownership are also receiving increasing consideration by groups such as the "Data Trusts" initiative [38] and others. Our goal is to support the growth of the DMLR ecosystem into a strong community with effective infrastructure that will advance machine learning science through next generation datasets. These datasets shall serve as a bridge to connect fundamental problems (such as food insecurity and climate impact) with fundamental ML research by providing the right datasets for the right problems.

3 Present: Convening the Community at ICML 2023

In order to charge this effort, the data-centric ML community came together on July 29, 2023, in Honolulu, Hawaii, for the inaugural DMLR workshop at the International Conference on Machine Learning (ICML) 2023. The DMLR workshop was a point of convergence

2. https://mlsys.org/
3. http://deem-workshop.org/
4. https://github.com/lancedb/lance
5. https://github.com/mlcommons/croissant

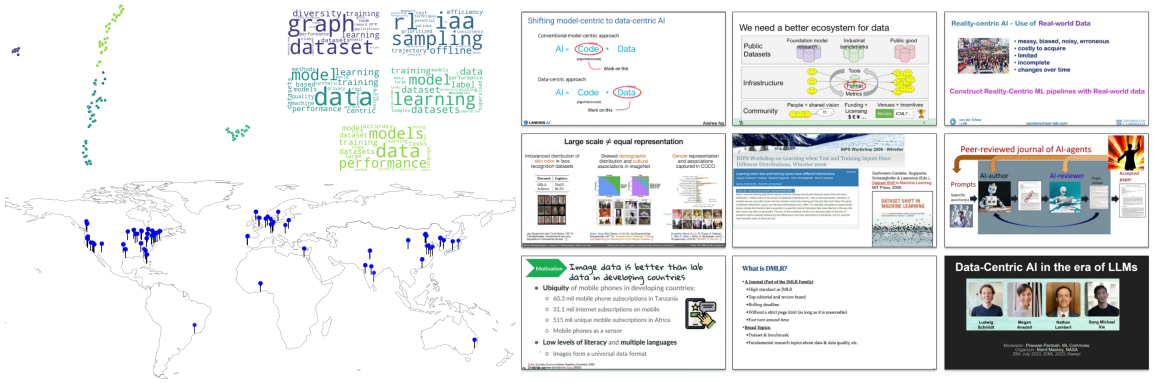

Figure 2: Themes and contributions from the community at the DMLR ICML 2023 workshop. **Top left:** LDA of accepted paper abstracts with `n_components` = 5. 2-d UMAP of LDA results which are 5-d corresponding to 5 components. Each dot represents an abstract, color coded by the most dominant topic identified by LDA. The topics identified by LDA are displayed alongside as top20 word clouds. **Bottom left:** A sample of the geographic coordinates of the institutions where authors of accepted works are based. It includes only those locations where the `geocode` API returns latitude and longitude information for fuzzy search on affiliation names (360 of 495 affiliations returned coordinates, note that not all 495 affiliations are unique). **Right:** Topics highlighted in the invited talk including prompt-based ML development (Andrew Ng), the DMLR ecosystem (Peter Mattson), reality-centric AI (Mihaela van der Schaar), bias in vision data (Olga Russakovsky and Vikram Ramaswamy), history of distribution shifts dating back to NeurIPS 2006 (Masashi Sugiyama), the AI research agent (Isabelle Guyon), nuances of data quality (Dina Machuve), the DMLR Journal (Ce Zhang) and data-centric LLMs (panel). Links to the full videos and slides of talks are available in Appendix C.

for previous activities including the Asilomar Datasets 2030 retreat, the Dataperf initiative[6], the NeurIPS 2021 data-centric AI workshop[7], the LAION community[8] and others. Invited speakers, panelists, poster presenters and attendees deliberated on the current state of data-centric machine learning and how we can advance the community and infrastructure towards the next generation of public machine learning datasets (see Figure 2 for a brief overview).

Community engagement Andrew Ng concluded his keynote with open questions aimed at fostering further research and development in data-centric AI workflows. Isabelle Guyon proposed a peer-reviewed journal contributed to by AI-agents, aiming to foster scholarly community engagement. Dina Machuve discussed the role of community in data collection for agriculture in East Africa. Olga Rusnaskovsky and Vikram V. Ramaswamy addressed social bias in machine learning, calling for community action. The panel expressed substantial enthusiasm for the DMLR Journal, indicating a strong community interest in advancing

6. `https://www.dataperf.org/`
7. `https://datacentricai.org/neurips21/`
8. `https://laion.ai/`

the field. Paper authors highlighted diverse challenges in community standards ranging from risk classification in driver telematics, the role of synthetic data in the scientific community, to the nuances of deep learning in neuroimaging and beyond.

Infrastructure Workshop contributions also illuminated the critical role of infrastructure in advancing data-centric machine learning. Andrew Ng emphasized the importance of rapid iteration cycles, facilitated by advancements in both theory and tools. Mihaela van der Schaar introduced tools like Data-IQ [39] for better data characterization. Peter Mattson and Praveen Paritosh discussed Croissant[9], a standardized dataset format, and DataPerf [17], an engine for refining datasets. Masashi Sugiyama added depth by discussing the complexities of machine learning models operating under distribution shifts. The panel, consisting of Ludwig Schmidt, Megan Ansdell, Nathan Lambert, and Sang Michael Xie, further emphasized that the development of systematic methods for constructing AI datasets is less advanced compared to model development but noted that tools and infrastructure are catching up. Poster presenters highlighted different aspects related to infrastructure such as quality control and streaming of distributed data.

Datasets The workshop participants also delved into the future of datasets in machine learning. Andrew Ng highlighted the growing relevance of small datasets and the practicality of few-shot learning techniques. Mihaela van der Schaar advocated for Reality-Centric AI [40]. Isabelle Guyon introduced AutoML+, a holistic system that includes data search, task definition, and preparation. Dina Machuve discussed the critical role of data in East African agriculture. The panel emphasized that data holds a central role in driving AI forward and highlighted the need for next-gen datasets to be more systematically constructed. Several papers also underscored the challenges and solutions in active learning, focusing on topics such as minimizing annotation cost and acquiring high-quality data for training discriminative models. Links to the full videos and slides of talks are available in Appendix C.

4 Future: Growing the DMLR Ecosystem

The field of machine learning is undergoing a profound transformation. While the past was characterized by the pursuit of innovative algorithms and architectures, the present and future pose growing data-centric questions. As large models become the norm and real-world efficacy becomes paramount, the emphasis is shifting towards the entire data lifecycle, from collection over storage and transformation to integration of results into other systems [41]. The importance of addressing societal issues through data further underscores this shift.

The role of the community in shaping the future of data-centric ML cannot be overstated. The recent DMLR workshop at ICML 2023 served as an inaugural meeting, igniting a spark for what is to come. A collective effort is required to create, enhance, and maintain public datasets. This involves establishing clear licensing protocols, technological standards, and fostering a culture of collaboration and shared, equitable ownership [42]. Interested contributors can already find several entry points, including

9. `https://github.com/mlcommons/croissant`

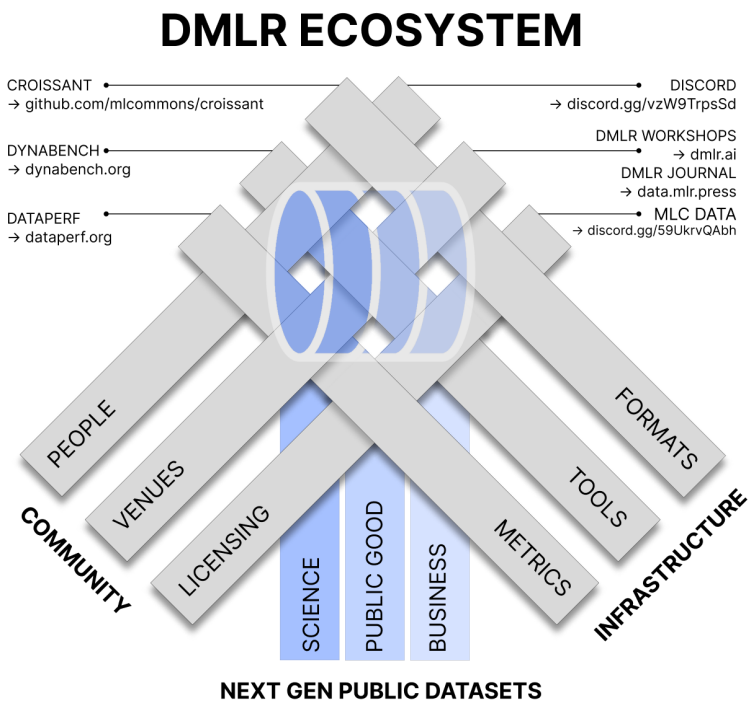

Figure 3: An overview of the DMLR ecosystem pillars and community projects.

- DMLR workshops or tracks at the main machine learning conferences such as NeurIPS Datasets and Benchmarks
  `https://dmlr.ai/`
- DMLR journal as an author, editor, reviewer
  `https://data.mlr.press/`
- Data provenance and governance initiatives such as
  `https://datatrusts.uk/` [38]
  `https://dataprovenance.org/` [42]
- Socials and informal research retreats such as Asilomar Datasets 2030
- Open source data-centric libraries such as
  `https://github.com/vanderschaarlab/datagnosis`

An infrastructure that supports the collaborative creation and enhancement of datasets is crucial. This infrastructure should champion the principles of open-source software, fostering a culture of shared responsibility and continuous improvement. The concept of "living datasets" emerges, emphasizing the dynamic nature of data [43; 44] and the importance of metrics [45; 46] and rich, flexible metadata in ensuring its relevance. Exemplar activities that continuously onboard input from contributors include, among others

- Croissant dataset format
  `https://github.com/mlcommons/croissant`

- Dynabench dynamic data collection and benchmarking platform
  `https://dynabench.org/`
- Dataperf, metrics for data-centric algorithm benchmarks
  `https://www.dataperf.org/home` [17]

Vibrant communities and innovative infrastructure will facilitate the future of machine learning datasets that cater to large models and real-world efficacy. These datasets should encapsulate the entire data lifecycle, ensuring they remain relevant and adaptable. They must be amenable enough to support the evolving research questions in machine learning. Furthermore, they should help address societal issues and allow analyses with respect representation and biases [41; 47]. The integrity of data forms the bedrock of reliable machine learning models. This involves addressing challenges related to noisy measurements, noisy labels and uncertainty [48; 49]. Ensuring the quality of data used for ML training and evaluation is paramount, as it directly influences the efficacy and reliability of the resulting models. New datasets, also called data++ [50] by some, thus should increasingly support the optimization of data itself [51–55] as part of the machine learning lifecycle. Ongoing initiatives that amalgamate these ingredients comprise, among others

- Machine Learning Common's datasets working group
  `https://mlcommons.org/en/groups/datasets/`
- UN's AI for Good SDG gateway
  `https://aiforgood.itu.int/gateway/`
- Independent research collectives such as LAION or EleutherAI
  `https://laion.ai/`, `https://www.eleuther.ai/`
- A diversity of open source benchmarking and evaluation repositories such as
  `https://github.com/erichson/SuperBench`,
  `https://wilds.stanford.edu/` [56],
  `https://github.com/hendrycks/robustness`[57],
  `https://github.com/basveeling/pcam` [58],
  `https://mlcommons.org/en/dollar-street/` [59],
  `https://github.com/modestyachts/ImageNetV2` [60],
  `https://github.com/inverse-scaling/prize` [61]

Investing in public datasets offers a plethora of benefits. It has the potential to accelerate innovation in the field of ML, reduce legal and ethical risks associated with data usage, and address pressing societal challenges. The emphasis is on creating datasets that not only advance the field of ML but also contribute positively to society at large by addressing real-world problems. The DMLR community is already expansive and, even more importantly, ongoing. We envision an ecosystem that strengthens these pillars and supports the growth and funding of new ideas. Whether you are a researcher, a practitioner, or an enthusiast, your insights and contributions to DMLR are the determinants of the data-centric machine learning future.

Continual learning without forgetting

With this editorial we aim to highlight critical developments in data-centric machine learning and provide an overview of entry points for contributions to different activities in the extended community. In a dynamic system, a snapshot like this editorial will always contain some approximation error. If you know of relevant resources that were omitted please do not be shy and reach out. We will be happy to update them.

References

[1] CL Wilson and MD Garris. Handprinted character database, nist special database 1. *NIST Technical Report and CDROM*, 1990.

[2] Michael D Garris, R Wilkinson, and Charles L Wilson. Methods for enhancing neural network handwritten character recognition. In *International Joint Conference on Neural Networks.*, volume 1, pages 695–700. Citeseer, 1991.

[3] MD Garris and RA Wilkinson. Handwritten segmented characters database. In *Technical Report Special Database*, volume 3. HWSC, National Institute of Standards and Technology, 1992.

[4] CL Wilson and MD Garris. Handprinted character database 3. Technical report, Technical report, National Institute of Standards and Technology, 1992.

[5] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[6] Chhavi Yadav and Léon Bottou. Cold case: The lost mnist digits. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019.

[7] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. "everyone wants to do the model work, not the data work": Data cascades in high-stakes ai. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.

[8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

[9] A Krizhevsky. Learning multiple layers of features from tiny images. *Master's thesis, University of Tront*, 2009.

[10] Yuri Breitbart, Andrew Deacon, Hans-Jörg Schek, Amit P Sheth, and Gerhard Weikum. Merging application-centric and data-centric approaches to support transaction-oriented multi-system workflows. *Sigmod record*, 22(3):23–30, 1993.

[11] Ronan Collobert, Samy Bengio, and Johnny Mariéthoz. Torch: a modular machine learning software library. 2002. URL `https://api.semanticscholar.org/CorpusID:15187647`.

[12] James Bergstra, Frédéric Bastien, Olivier Breuleux, Pascal Lamblin, Razvan Pascanu, Olivier Delalleau, Guillaume Desjardins, David Warde-Farley, Ian J. Goodfellow, Arnaud Bergeron, and Yoshua Bengio. Theano: Deep learning on gpus with python. 2012. URL `https://api.semanticscholar.org/CorpusID:62497190`.

[13] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

[14] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning, 2016.

[15] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

[16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019.

[17] Mark Mazumder, Colby Banbury, Xiaozhe Yao, Bojan Karlaš, William Gaviria Rojas, Sudnya Diamos, Greg Diamos, Lynn He, Alicia Parrish, Hannah Rose Kirk, Jessica Quaye, Charvi Rastogi, Douwe Kiela, David Jurado, David Kanter, Rafael Mosquera, Juan Ciro, Lora Aroyo, Bilge Acun, Lingjiao Chen, Mehul Smriti Raje, Max Bartolo, Sabri Eyuboglu, Amirata Ghorbani, Emmett Goodman, Oana Inel, Tariq Kane, Christine R. Kirkpatrick, Tzu-Sheng Kuo, Jonas Mueller, Tristan Thrush, Joaquin Vanschoren, Margaret Warren, Adina Williams, Serena Yeung, Newsha Ardalani, Praveen Paritosh, Lilith Bat-Leah, Ce Zhang, James Zou, Carole-Jean Wu, Cody Coleman, Andrew Ng, Peter Mattson, and Vijay Janapa Reddi. Dataperf: Benchmarks for data-centric ai development, 2023.

[18] Francis Ysidro Edgeworth. On the probable errors of frequency-constants. *Journal of the Royal Statistical Society*, 71(2):381–397, 1908.

[19] R. A. Fisher and Edward John Russell. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 222(594-604):309–368, 1922. doi: 10.1098/rsta.1922.0009. URL `https://royalsocietypublishing.org/doi/abs/10.1098/rsta.1922.0009`.

[20] N.Z. Tishby and R.D. Levine. Surprisal analysis derived from a variational principle for mechanical systems. *Chemical Physics Letters*, 98(4):310–314, 1983. ISSN

0009-2614. doi: https://doi.org/10.1016/0009-2614(83)80213-9. URL https://www.sciencedirect.com/science/article/pii/0009261483802139.

[21] Les Atlas, David Cohn, and Richard Ladner. Training connectionist networks with queries and selective sampling. *Advances in neural information processing systems*, 2, 1989.

[22] Sebastian B Thrun and Knut Möller. Active exploration in dynamic environments. *Advances in neural information processing systems*, 4, 1991.

[23] David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.

[24] David Cohn, Zoubin Ghahramani, and Michael Jordan. Active learning with statistical models. In G. Tesauro, D. Touretzky, and T. Leen, editors, *Advances in Neural Information Processing Systems*, volume 7. MIT Press, 1994. URL https://proceedings.neurips.cc/paper_files/paper/1994/file/7f975a56c761db6506eca0b37ce6ec87-Paper.pdf.

[25] Jan Storck, Sepp Hochreiter, Jürgen Schmidhuber, et al. Reinforcement driven information acquisition in non-deterministic environments. In *Proceedings of the international conference on artificial neural networks, Paris*, volume 2, pages 159–164, 1995.

[26] Isabelle Guyon, Gavin C. Cawley, Gideon Dror, and Vincent Lemaire. Results of the active learning challenge. In Isabelle Guyon, Gavin Cawley, Gideon Dror, Vincent Lemaire, and Alexander Statnikov, editors, *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*, volume 16 of *Proceedings of Machine Learning Research*, pages 19–45, Sardinia, Italy, 16 May 2011. PMLR. URL https://proceedings.mlr.press/v16/guyon11a.html.

[27] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III au2, and Kate Crawford. Datasheets for datasets, 2021.

[28] Ruiyang Zhao, Burhaneddin Yaman, Yuxin Zhang, Russell Stewart, Austin Dixon, Florian Knoll, Zhengnan Huang, Yvonne W. Lui, Michael S. Hansen, and Matthew P. Lungren. fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil mri data, 2021.

[29] Judy Wawira Gichoya, Imon Banerjee, Ananth Reddy Bhimireddy, John L Burns, Leo Anthony Celi, Li-Ching Chen, Ramon Correa, Natalie Dullerud, Marzyeh Ghassemi, Shih-Cheng Huang, Po-Chih Kuo, Matthew P Lungren, Lyle J Palmer, Brandon J Price, Saptarshi Purkayastha, Ayis T Pyrros, Lauren Oakden-Rayner, Chima Okechukwu, Laleh Seyyed-Kalantari, Hari Trivedi, Ryan Wang, Zachary Zaiman, and Haoran Zhang. AI recognition of patient race in medical imaging: a modelling study. *The Lancet Digital Health*, 4(6):e406–e414, jun 2022. doi: 10.1016/s2589-7500(22)00063-2. URL https://doi.org/10.1016%2Fs2589-7500%2822%2900063-2.

[30] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5389–5400. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/recht19a.html.

[31] Dongyoon Han, Junsuk Choe, Seonghyeok Chun, John Joon Young Chung, Minsuk Chang, Sangdoo Yun, Jean Y. Song, and Seong Joon Oh. Neglected free lunch – learning image classifiers using annotation byproducts. In *International Conference on Computer Vision (ICCV)*, 2023.

[32] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.

[33] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

[34] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *arXiv preprint arXiv:2305.01569*, 2023.

[35] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagenet hierarchy. In *Conference on Fairness, Accountability, and Transparency*, 2020. doi: 10.1145/3351095.3375709.

[36] Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky. A study of face obfuscation in imagenet. In *International Conference on Machine Learning (ICML)*, 2022.

[37] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.

[38] Sylvie Delacroix and Neil D Lawrence. Bottom-up data Trusts: disturbing the 'one size fits all' approach to data governance. *International Data Privacy Law*, 9(4):236–252, 10 2019. ISSN 2044-3994. doi: 10.1093/idpl/ipz014. URL https://doi.org/10.1093/idpl/ipz014.

[39] Nabeel Seedat, Jonathan Crabbé, Ioana Bica, and Mihaela van der Schaar. Data-IQ: Characterizing subgroups with heterogeneous outcomes in tabular data. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=qC2BwvfaNdd.

[40] Mihaela van der Schaar and Andrew Rashbass. The case for reality-centric ai // van der schaar lab. https://www.vanderschaar-lab.com/the-case-for-reality-centric-ai/. (Accessed on 10/18/2023).

[41] Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, L Fei-Fei, Matei Zaharia, Ce Zhang, and James Zou. Advances, challenges and opportunities in creating data for trustworthy ai. *Nature Machine Intelligence*, 4(8):669–677, 2022.

[42] Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, Xinyi Wu, Enrico Shippole, Kurt Bollacker, Tongshuang Wu, Luis Villa, Sandy Pentland, Deb Roy, and Sara Hooker. The data provenance initiative: A large scale audit of dataset licensing and attribution in ai, 2023.

[43] Luis Oala. *Metrological machine learning (2ML)*. 1 edition, 2023. URL https://metrological.ml.

[44] Ian R. McKenzie, Alexander Lyzhov, Michael Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Aaron Kirtland, Alexis Ross, Alisa Liu, Andrew Gritsevskiy, Daniel Wurgaft, Derik Kauffman, Gabriel Recchia, Jiacheng Liu, Joe Cavanagh, Max Weiss, Sicong Huang, The Floating Droid, Tom Tseng, Tomasz Korbak, Xudong Shen, Yuhui Zhang, Zhengping Zhou, Najoung Kim, Samuel R. Bowman, and Ethan Perez. Inverse scaling: When bigger isn't better, 2023.

[45] Dan Friedman and Adji Bousso Dieng. The vendi score: A diversity evaluation metric for machine learning, 2023.

[46] Amey Pasarkar and Adji Bousso Dieng. Cousins of the vendi score: A family of similarity-based diversity metrics for science and machine learning, 2023.

[47] Vikram V. Ramaswamy, Sing Yu Lin, Dora Zhao, Aaron B. Adcock, Laurens van der Maaten, Deepti Ghadiyaram, and Olga Russakovsky. GeoDE: a geographically diverse evaluation dataset for object recognition. In *NeurIPS Datasets and Benchmarks Track*, 2023.

[48] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the imagenet hierarchy. *CoRR*, abs/1912.07726, 2019. URL http://arxiv.org/abs/1912.07726.

[49] Curtis G Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. URL https://openreview.net/forum?id=XccDXrDNLek.

[50] Phillip Isola. [gcv @ cvpr23] generative models as data++. https://www.youtube.com/watch?v=YuRAeQsTSo8, 2023. (Accessed on 10/18/2023).

[51] Manel Baradad Jurjo, Jonas Wulff, Tongzhou Wang, Phillip Isola, and Antonio Torralba. Learning to see by looking at noise. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in*

Neural Information Processing Systems, volume 34, pages 2556–2569. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/14f2ebeab937ca128186e7ba876faef9-Paper.pdf.

[52] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 20596–20607. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/ac56f8fe9eea3e4a365f29f0f1957c55-Paper.pdf.

[53] Nabeel Seedat, Fergus Imrie, and Mihaela van der Schaar. Dc-check: A data-centric ai checklist to guide the development of reliable machine learning systems, 2022.

[54] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=qhAeZjs7dCL.

[55] Luis Oala, Marco Aversa, Gabriel Nobis, Kurt Willis, Yoan Neuenschwander, Michèle Buck, Christian Matek, Jerome Extermann, Enrico Pomarico, Wojciech Samek, Roderick Murray-Smith, Christoph Clausen, and Bruno Sanguinetti. Data models for dataset drift controls in machine learning with optical images. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=I4IkGmgFJz.

[56] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. Wilds: A benchmark of in-the-wild distribution shifts, 2021.

[57] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.

[58] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant CNNs for digital pathology. June 2018.

[59] William A Gaviria Rojas, Sudnya Diamos, Keertan Ranjan Kini, David Kanter, Vijay Janapa Reddi, and Cody Coleman. The dollar street dataset: Images representing the geographic and socioeconomic diversity of the world. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL https://openreview.net/forum?id=qnfYsave0U4.

[60] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do cifar-10 classifiers generalize to cifar-10? 2018. https://arxiv.org/abs/1806.00451.

[61] Ian McKenzie, Alexander Lyzhov, Alicia Parrish, Ameya Prabhu, Aaron Mueller, Najoung Kim, Sam Bowman, and Ethan Perez. The inverse scaling prize, 2022. URL `https://github.com/inverse-scaling/prize`.

Appendix A. The people behind the DMLR program at ICML 2023

Next to the organizers, speakers and attendees, the DMLR community is made up of its reviewers and submitting authors. For the first DMLR meeting at ICML 2023 (`https://dmlr.ai/`) we are grateful to the following people for volunteering their time and expertise.

A.1 Program committee

Ziniu Li (The Chinese University of Hong Kong, Shenzhen), Zhixin Huang (University of Kassel), Zhaowei Zhu (University of California, Santa Cruz), Yue Yu (Georgia Institute of Technology), Yue Xing (Michigan State University), Yuanshun Yao (ByteDance), Yoav Wald (Johns Hopkins), Yixin Liu (Monash University), Yilin Zhang (Meta), Yi-Fan Zhang (NLPR, China), Yang Liu (UC Santa Cruz), Xinhui Li (Georgia Institute of Technology), Xianling Zhang (Ford Motor Company), William Gaviria Rojas (Coactive AI), Usama Muneeb (University of Illinois Chicago), Tzu-Sheng Kuo (Carnegie Mellon University), Tom Viering (Delft University of Technology, Netherlands), Thao Nguyen (University of Washington), Sumedh Datar (UTA), Sigrid Passano Hellan (University of Edinburgh), Siddharth Joshi (UCLA), Si Chen (Virginia Tech), Shin'ya Yamaguchi (NTT / Kyoto University), Sebastian Schelter (University of Amsterdam), Roger Waleffe (University of Wisconsin-Madison), Rasool Fakoor (AWS), Puja Trivedi (University of Michigan), Praveen Paritosh (Google), Peter Mattson (Google), Paolo Climaco (Universitat Bonn), Oliver Lenz (Universiteit Gent), Nauman Ahad (Georgia Institute of Technology), Muhammed Razzak (University of Oxford), Min Du (Palo Alto Networks), Megan Richards (Meta), Mayee Chen (Stanford University), Manil Maskey (NASA MSFC), Madelon Hulsebos (University of Amsterdam), Luis Oala (Dotphoton AG), Linxin Song (Waseda University), Linus Ericsson (University of Edinburgh), Lilith Bat-Leah (N/A), Liangchen Luo (Google), Li Jiang (Tsinghua University), Lenora Gray (Redgrave Data), Kurt Bollacker (The Long Now Foundation), Karthick Gunasekaran (Researcher), Julian Bitterwolf (University of Tubingen), Jinyi Liu (Tianjin University), Jieyu Zhang (University of Washington), Jialu Wang (University of California, Santa Cruz), Jiaheng Wei (UCSC), Jiachen Wang (Princeton University), Jeyeon Eo (Soongsil University), Jerone Andrews (Sony AI), Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), Jarne Van den Herrewegen (Oqton / Ghent University), Jan Van Rijn (Leiden University), Ian Beaver (Verint Systems Inc), Huaizheng Zhang (BreezeML), Himchan Jeong (Simon Fraser University), Hidetomo Sakaino (Weathernews Inc.), Harit Vishwakarma (University of Wisconsin Madison), Hao Cheng (University of California, Santa Cruz), Hang Zhou (UC Davis), Guozheng Ma (Tsinghua University), Gregory Yauney (Cornell University), Feiyang Kang (Virginia Tech), Fangyi Chen (Carnegie Mellon University), Dionysis Manousakas (Amazon), Diego Botache (University of Kassel), Danilo Brajovic (Fraunhofer), Daniel Galvez (NVIDIA), Chanjun Park (Upstage), Beverly Quon (University of California, Irvine), Andre Carreiro (Fraunhofer Portugal AICOS), Amro Abbas (Meta), Ali Hakimi Parizi (Thomson Reuters), Alexander Li (Carnegie Mellon University)

A.2 Authors

Hashan Peiris (Simon Fraser University), Himchan Jeong (Simon Fraser University), Jae Kwang Kim (Iowa State University), Gantavya Bhatt (University of Washington, Seattle), Arnav Das (University of Washington), Megh M Bhalerao (University of Washington), Rui Yang (Memorial Sloan Kettering Cancer Center), Vianne R Gao (Weill Medical College), Jeff Bilmes (UW), Linxin Song (Waseda University), Jieyu Zhang (University of Washington), Xiaotian Lu (Kyoto University), Tianyi Zhou (University of Maryland, College Park), Yue Xing (Michigan State University), Ashutosh Pandey (Meta Platforms), David Yan (Meta Platforms), Fei Wu (Meta), Michael Fronda (Meta Platforms), Pamela Bhattacharya (Meta Platforms), Yongchao Zhou (University of Toronto), Hshmat U Sahak (University of Toronto), Jimmy Ba (University of Toronto), Eujeong Choi (Upstage), Chanjun Park (Upstage), NamHyeok Kim (Upstage), Damrin Kim (Konkuk University), Harksoo Kim (Konkuk University), Sang Michael Xie (Stanford University), Hieu Pham (Google), Xuanyi Dong (University of Technology Sydney), Nan Du (Google Brain), Hanxiao Liu (Google Brain), Yifeng Lu (Google Brain), Percy Liang (Stanford University), Quoc Le (Google Brain), Tengyu Ma (Stanford), Adams Wei Yu (Google Brain), Bohan Wang (University of Science and Technology of China), zhengyu hu (NA), Pang We Koh (University of Washington), Alexander J Ratner (University of Washington), Jifan Zhang (University of Wisconsin), Shuai Shao (Meta), Saurabh Verma (Meta), Robert Nowak (University of Wisconsin, Madison), Ziniu Li (The Chinese University of Hong Kong, Shenzhen), Tian Xu (Nanjing University), Zeyu Qin (HKUST), Yang Yu (Nanjing University), Zhiquan Luo (The Chinese University of Hong Kong, Shenzhen and Shenzhen Research Institute of Big Data), Seonmin Koo (Korea University), Seolhwa Lee (University of Copenhagen), Jaehyung Seo (Korea University), Sugyeong Eo (Korea University), Hyeonseok Moon (Korea University), Heuiseok Lim (Korea University), Piotr Przemielewski (Jagiellonian University), Witold Wydmański (Jagiellonian University), Marek Śmieja (Jagiellonian University), Hang Zhou (UC Davis), Jonas Mueller (Cleanlab), Mayank Kumar (Cleanlab), Jane-Ling Wang (UC Davis), Jing Lei (Carnegie Mellon University), Saad A Almohaimeed (University of Central Florida), Saleh Almohaimeed (University of Central Florida), Ashfaq Ali Shafin (Florida International University), Bogdan Carbunar (Florida International University), Ladislau Boloni (University of Central Florida), Yuanshun Yao (ByteDance), Yang Liu (UC Santa Cruz), Harit Vishwakarma (University of Wisconsin Madison), Heguang Lin (University of Wisconsin-Madison), Frederic Sala (University of Wisconsin-Madison), Ramya Korlakai Vinayak (University of Wisconsin-Madison), Jesse E Cummings (MIT), Elías Snorrason (Cleanlab), Seungjun Lee (Korea University), Stefan Grafberger (University of Amsterdam), Bojan Karlaš (Harvard University), Paul Groth (University of Amsterdam), Sebastian Schelter (University of Amsterdam), Huaizheng Zhang (BreezeML), Yizheng Huang (BreezeML), Yuanming Li (Independent Researcher), Jaeseung Heo (POSTECH), Seungbeom Lee (POSTECH), Sungsoo Ahn (POSTECH), Dongwoo Kim (POSTECH), Patrik Okanovic (ETH Zurich), Roger Waleffe (University of Wisconsin-Madison), Vasilis Mageirakos (ETH Zurich), Konstantinos Nikolakakis (Yale University), Amin Karbasi (Yale), Dionysios Kalogerias (Yale University), Nezihe Merve Gürel (ETH Zürich), Theodoros Rekatsinas (ETH Zurich), Si Chen (Virginia Tech), Feiyang Kang (Virginia Tech), Nikhil Abhyankar (Virginia Tech), Ming Jin

(Virginia Tech), Ruoxi Jia (Virginia Tech), Joshua L Vendrow (MIT), Saachi Jain (MIT), Logan Engstrom (MIT), Aleksander Madry (MIT), Hoang Anh Just (Virginia Tech), Anit Kumar Sahu (Amazon Alexa AI), Jinsung Kim (Korea University), Min Du (Palo Alto Networks), Nesime Tatbul (Intel Labs and MIT), Brian Rivers (Intel), Akhilesh Kumar Gupta (University of Pennsylvania), Lucas Hu (Palo Alto Networks), Wei Wang (Palo Alto Networks), Ryan C Marcus (MIT), Shengtian Zhou (Snap), Insup Lee (University of Pennsylvania), Justin Gottschlich (Merly and Stanford University), Paolo Climaco (Institut für Numerische Simulation, Universität Bonn), Jochen Garcke (University Bonn), Nathan Vaska (MIT Lincoln Laboratories), Victoria Helus (MIT Lincoln Laboratory), Natalie Abreu (MIT Lincoln Laboratory), Dahyun Jung (Korea University), Jaewook Lee (Korea University), Yue Yu (Georgia Institute of Technology), Yuchen Zhuang (Georgia Institute of Technology), Yu Meng (University of Illinois Urbana-Champaign), Ranjay Krishna (University of Washington), Jiaming Shen (Google Research), Chao Zhang (Georgia Institute of Technology), Jerone T A Andrews (Sony AI), Dora Zhao (Sony AI), William Thong (Sony AI), Apostolos Modas (Sony), Orestis Papakyriakopoulos (Sony AI), Alice Xiang (Sony AI), Lei Shu (Google), Liangchen Luo (Google), Jayakumar Hoskere (Google), Yun Zhu (Google), Yinxiao Liu (Google), Simon Tong (Google), Jindong Chen (Google), Lei Meng (Google), Yongchan Kwon (Columbia University), James Zou (Stanford University), Shivangana Rawat (Indian Institute of Technology, Hyderabad), Chaitanya Devaguptapu (Fujitsu Research), Vineeth Balasubramanian (Indian Institute of Technology Hyderabad), Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), Vivek Narayanaswamy (Lawrence Livermore National Laboratory), Puja Trivedi (University of Michigan), Rushil Anirudh (Lawrence Livermore National Laboratory), Shreyas Krishnaswamy (University of California, Berkeley), Lisa Dunlap (UC Berkeley), Lingjiao Chen (University of Wisconsin-Madison), Matei Zaharia (Stanford and Databricks), Joey Gonzalez (Berkeley), Hao Cheng (University of California, Santa Cruz), Qingsong Wen (Alibaba DAMO Academy), Liang Sun (Alibaba Group), Mayee Chen (Stanford University), Nicholas Roberts (University of Wisconsin-Madison), Kush Bhatia (Stanford University), Jue WANG (Zhejiang University), Ce Zhang (ETH), Christopher Re (Stanford University), Andre V Carreiro (Fraunhofer Portugal AICOS), Mariana Pinto (Faculty of Science and Technology, Nova University of Lisbon), Pedro S Madeira (Fraunhofer Portugal AICOS), Alberto Lopez (Imprensa Nacional - Casa da Moeda), Hugo Gamboa (LIBPhys, Faculdade de Ciências e Tecnologia, Universidade Noval de Lisboa), M. Eren Akbiyik (ETH Zurich), Florian Grötschla (ETH Zürich), Beni Egressy (ETH Zurich), Roger Wattenhofer (ETH Zurich), Oliver U Lenz (Universiteit Gent), Daniel Peralta (Ghent University), Chris Cornelis (Ghent University), Aabha Pingle (Pune Institute of Computer Technology), Aditya Vyawahare (Pune Institute of Computer Technology), Isha Joshi (Pune Institute of Computer Technology), Rahul Tangsali (SCTR's Pune Institute of Computer Technology), Raviraj Joshi (Indian Institute of Technology Madras), Jarne Van den Herrewegen (Oqton / Ghent University), Tom Tourwé (Oqton), Francis Wyffels (Ghent University), Julian Bitterwolf (University of Tübingen), Maximilian Mueller (University of Tübingen), Matthias Hein (University of Tübingen), Darlington Akogo (minoHealth), Issah A Samori (minoHealth AI Labs), Cyril S K Akafia (minoHealth AI Labs), Harriet Dede Fiagbor (minoHealth AI Labs), Andrews A Kangah (KaraAgro AI Labs), Donald Donald (KaraAgro), Kwabena Fuachie (Kara Agro AI), Luis Oala (Dotphoton AG), Li Jiang (Tsinghua University), Sijie Chen (Fudan Uni-

versity), Jielin Qiu (Carnegie Mellon University), Haoran Xu (JD Technology), Victor Chan (TBSI), DING ZHAO (Carnegie Mellon University), Sigrid Passano Hellan (University of Edinburgh), Christopher Lucas (University of Edinburgh), Nigel Goddard (University of Edinburgh), Dionysis Manousakas (Amazon), Sergul Aydore (Amazon), Nauman Ahad (Georgia Institute of Technology), Namrata Nadagouda (Georgia Institute of Technology), Eva L Dyer (Georgia Tech), Mark Davenport (Georgia Institute of Technology), Rafael Mosquera Gómez (MLCommons), Julian Eusse (MLCommons), Juan Manual Ciro (Factored), Daniel Galvez (NVIDIA), Ryan Hileman (Talon Voice), Kurt Bollacker (The Long Now Foundation), David Kanter (MLCommons), Siddharth Joshi (UCLA), Baharan Mirzasoleiman (UCLA), Jeffrey Li (University of Washington), Ludwig Schmidt (University of Washington), Vedang Lad (MIT), Ammar Sherif (Nile University), Abubakar Abid (Hugging Face), Mustafa Elattar (Nile University), Mohamed ElHelw (Nile University), Ulyana Tkachenko (Cleanlab), Aditya Thyagarajan (CleanLab), Alycia Y Lee (Stanford University), Brando Miranda (Stanford University), Sanmi Koyejo (Stanford University), Xinhui Li (Georgia Institute of Technology), Alex Fedorov (Georgia Institute of Technology), Mrinal Mathur (Georgia State University), Anees Abrol (TReNDS), Gregory Kiar (Child Mind Institute), Sergey Plis (Georgia State University), Vince Calhoun (TReNDS), Patrick Yu (University of Illinois Urbana-Champaign), Saumya Goyal (Stanford University), Yu-Xiong Wang (University of Illinois at Urbana-Champaign), Beverly A Quon (University of California, Irvine), Jean-Luc Gaudiot (University of California, Mark Heimann (Lawrence Livermore), Danai Koutra (U Michigan), Yifang Chen (University of Washington), Gregory H Canal (University of Wisconsin-Madison), Stephen O Mussmann (University of Washington), Yinglun Zhu (University of Wisconsin-Madison), Simon Du (University of Washington), Kevin Jamieson (U Washington), Alexander C Li (Carnegie Mellon University), Ellis L Brown (Carnegie Mellon University), Alexei A Efros (UC Berkeley), Deepak Pathak (Carnegie Mellon University), Dingshuo Chen (University of Chinese Academy of Sciences), Yanqiao ZHU (University of California, Los Angeles), Yuanqi Du (Cornell University), Zhixun Li (The Chinese University of Hong Kong), Qiang Liu (Institute of Automation, Chinese Academy of Sciences), Shu Wu (NLPR, China), Liang Wang (NLPR, Joel Niklaus (University of Bern), Veton Matoshi (Bern University of Applied Sciences), Matthias Stürmer (University of Bern), Ilias Chalkidis (University of Copenhagen), Daniel Ho (Stanford Law), Siddarth Ramesh (Adobe), Surgan Jandial (MDSR Labs, Adobe), Gauri Gupta (MIT), Piyush Gupta (Adobe Systems India Pvt Ltd), Balaji Krishnamurthy (), Kushal Tirumala (FAIR), Daniel Simig (Meta AI), Armen Aghajanyan (FAIR), Ari S Morcos (Facebook AI Research (FAIR)), Yonghyun Kwon (Iowa State University), Rohith Peddi (The University of Texas at Dallas), Shivvrat Arya (The University of Texas at Dallas), Bharath Challa (The University of Texas at Dallas), Likhitha Pallapothula (University of Texas at Dallas), AKSHAY VYAS (University of Texas at Dallas), Qifan Zhang (The University of Texas at Dallas), Jikai Wang (University of Texas at Dallas), Vasundhara Komaragiri (UT Dallas), Eric Ragan (University of Florida), Nicholas Ruozzi (UT Dallas), Yu Xiang (The University of Texas at Dallas), Vibhav Gogate (UT Dallas), Shin'ya Yamaguchi (NTT / Kyoto University), Daiki Chijiwa (NTT), Sekitoshi Kanai (NTT), Atsutoshi Kumagai (NTT Computer and Data Science Laboratories), Hisashi Kashima (Kyoto University), Jiaheng Wei (UCSC), Zhaowei Zhu (University of California, Tianyi Luo (Amazon), Ehsan Amid (Google Brain), Abhishek Kumar (Google Brain), Muhammed T Razzak (University of Oxford), Anthony

Ortiz (Microsoft), Caleb Robinson (Microsoft AI for Good Research Lab), Fangyi Chen (Carnegie Mellon University), Han Zhang (CMU), Hao Chen (Carnegie Mellon University), Kai Hu (Carnegie Mellon University), Jiachen Dou (Carnegie Mellon University), zaiwang li (pitt), Chenchen Zhu (Meta), Marios Savvides (Carnegie Mellon University), A. Feder Cooper (Cornell University), Wentao Guo (Cornell University), Duc Khiem Pham (Cornell University), Tiancheng Yuan (Cornell University), Charlie F Ruan (Cornell University), Yucheng Lu (Cornell University), Christopher De Sa (Cornell University), Rasool Fakoor (AWS), Zachary Lipton (Carnegie Mellon University), Pratik A Chaudhari (University of Pennsylvania), Alex J Smola (Amazon), Mark Vero (ETH Zurich), Mislav Balunovic (ETH Zurich), Martin Vechev (ETH Zurich), Jinyi Liu (Tianjin University), Yi Ma (Tianjin University), Jianye Hao (Tianjin University), Yujing Hu (NetEase Fuxi AI Lab), Yan Zheng (Tianjin University), Tangjie Lv (NetEase Fuxi AI Lab), Changjie Fan (NetEase Fuxi AI Lab), Gregory Yauney (Cornell University), Emily Reif (Google), David Mimno (Cornell University), Hailey Joren (UC San Diego), Chirag Nagpal (Carnegie Mellon University), Katherine Heller (Google), Berk Ustun (UCSD), Alex Oesterling (Harvard University), Jiaqi Ma (University of Illinois Urbana-Champaign), Flavio Calmon (Harvard University), Himabindu Lakkaraju (Harvard), Luísa B Shimabucoro (Universidade de São Paulo), Timothy Hospedales (Edinburgh University), Henry Gouk (University of Edinburgh), Pratyush Maini (IIT Delhi), Sachin Goyal (Carnegie Mellon University), Zico Kolter (Carnegie Mellon University), Aditi Raghunathan (Carnegie Mellon University), Aaditya Naik (University of Pennsylvania), Yinjun Wu (University of Pennsylvania), Mayur Naik (University of Pennsylvania), Eric Wong (University of Pennsylvania), Karthick Gunasekaran (Researcher), Sang Keun Choe (Carnegie Mellon University), Sanket Vaibhav Mehta (Carnegie Mellon University), Hwijeen Ahn (Carnegie Mellon University), Willie Neiswanger (Stanford University), Pengtao Xie (UC San Diego), Emma Strubell (Carnegie Mellon University), Eric Xing (MBZUAI, CMU, and Petuum Inc.), Guozheng Ma (Tsinghua University), Linrui Zhang (Tsinghua University), Haoyu Wang (Tsinghua University), Lu Li (Tsinghua University), Zilin Wang (Tsinghua University), Zhen Wang (The University of Sydney), Li Shen (JD Explore Academy), Xueqian Wang (Tsinghua University), Dacheng Tao (The University of Sydney), Yi-Fan Zhang (NLPR, Xue Wang (Alibaba DAMO Academy), Weiqi Chen (Alibaba Group), Zhang Zhang (Institute of Automation, Rong Jin (Twitter), Tieniu Tan (NLPR, Jiachen T. Wang (Princeton University), Yuqing Zhu (UC Santa Barbara), Yu-Xiang Wang (UC Santa Barbara), Prateek Mittal (Princeton University), Ching-Yun Ko (MIT), Pin-Yu Chen (IBM Research), Payel Das (IBM Research), Yung-Sung Chuang (MIT), Luca Daniel (Massachusetts Institute of Technology), Young In Kim (Purdue University), Pratiksha Agrawal (Purdue University), Johannes Royset (Naval Postgraduate School), Rajiv Khanna (Purdue University), Megan Richards (Meta), Diane Bouchacourt (Meta), Mark Ibrahim (Meta), Polina Kirichenko (New York University), Chiyuan Zhang (MIT), Linus Ericsson (University of Edinburgh), Newsha Ardalani (Meta AI (FAIR)), Mostafa Elhoushi (Meta), Carole-Jean Wu (Meta AI), Jacob Buckman (Mila), Kshitij Gupta (Mila), Ethan Caballero (Mila), Rishabh Agarwal (Google Research, Brain Team), Marc G. Bellemare (Google Brain), Avni Kothari (UC San Diego), Lily Weng (UCSD), Bogdan Kulynych (EPFL), Yoav Wald (Johns Hopkins), Suchi Saria (Johns Hopkins University), Hanyang Jiang (Georgia Institute of Technology), Yao Xie (Georgia Tech), Ellen Zegura (Georgia Tech), Elizabeth Belding (University of California, Santa Barbara), Shaowu Yuchi

(Georgia Institute of Technology), Kaize Ding (Arizona State University), Yancheng Wang (Arizona State University), Huan Liu (Arizona State University), Jeyeon Eo (Soongsil University), Dongsu Lee (Soongsil University), Minhae Kwon (Soongsil University), Thao T Nguyen (University of Washington), Samir Gadre (Columbia University), Gabriel Ilharco (University of Washington), Sewoong Oh (University of Washington), Kimia Hamidieh (University of Toronto, Vector Institute), Haoran Zhang (MIT), Thomas Hartvigsen (MIT), Marzyeh Ghassemi (University of Toronto, Amro Abbas (Meta), Surya Ganguli (Stanford University), Hidetomo Sakaino (Weathernews Inc.)

Appendix B. Full list of accepted papers

The full list of accepted papers is available at `https://dmlr.ai/accepted/`.

Appendix C. Links to recorded talks and slides from DMLR

In random order:

Masashi Sugiyama
Coping with Wild Distribution Shifts: Continuous Shift, Joint Shift, and Beyond
https://slideslive.com/39006435/coping-with-wild-distribution-shifts-continuous-shift-joint-shift-and-beyond?ref=folder-122509

Ce Zhang
DMLR: Journal of Data-centric Machine Learning Research
https://slideslive.com/39006439/dmlr-journal-of-datacentric-machine-learning-research?ref=folder-122509

Dina Machuve
Data for Agriculture: Challenges and Opportunities in East Africa
https://slideslive.com/39006438/data-for-agriculture-challenges-and-opportunities-in-east-africa?ref=folder-122509

Peter Mattson
Data-centric Ecosystem: Croissant and Dataperf
https://slideslive.com/39006431/datacentric-ecosystem-croissant-and-dataperf?ref=folder-122509

Olga Russakovsky and Vikram Ramaswamy
Data-centric Machine Learning: Tackling social bias in computer vision datasets
https://slideslive.com/39006434/datacentric-machine-learning-tackling-social-bias-in-computer-vision-datasets?ref=folder-122509

Andrew Ng
Fast prompt-based ML development and data-centric AI
https://slideslive.com/39006430/fast-promptbased-ml-development-and-datacentric-ai?ref=folder-122509

Ludwig Schmidt, Megan Ansdell, Nathan Lambert, Sang Michael Xie, Praveen Paritosh, Manil Maskey
Panel Discussion
https://slideslive.com/39006440/panel-discussion?ref=folder-122509

Mihaela van der Schaar
Reality-Centric AI
https://slideslive.com/39006433/realitycentric-ai?ref=folder-122509

Isabelle Guyon
Towards Data-Centric AutoML
https://slideslive.com/39006437/towards-datacentric-automl?ref=folder-122509