

Potential Impact of Data-Centric AI on Society

Sushant Kumar, Ritesh Sharma, and Vishakha Singh

Department of Computer Science and Engineering
Indian Institute of Technology (IIT BHU)
Varanasi 221005, India

Shrikant Tiwari

Department of Computer Science and Engineering (CSE)
Institute of Engineering and Rural Technology (IERT)
Prayagraj 211002, India

Sanjay Kumar Singh

Department of Computer Science and Engineering
Indian Institute of Technology (IIT BHU)
Varanasi 221005, India

Sumit Datta

School of Electronic Systems and Automation
IIITM-Kerala/Digital University Kerala
Thiruvananthapuram 695317, India

■ **DATA-CENTRIC ARTIFICIAL** intelligence (AI) (DCAI) has the potential to bring significant benefits to society; however, it also poses significant challenges and potential risks. It is crucial to approach the development and deployment of DCAI systems with caution, taking into account the potential societal impacts and working to mitigate any negative effects. DCAI technology is now an essential part of operations for many of the world's largest software and hardware industries. These industries offer a range of AI and machine-learning (ML) services, tools, and platforms to society to help businesses process and analyze data. By leveraging data, these industries are able to drive innovation, optimize their operations, and gain a competitive advantage in the market. From personalized recommendations to optimized manufacturing processes, data analytics and ML algorithms are being




Figure 1. Role of DCAI for the development of society.

used to improve the overall customer experience, increase efficiency, and identify new opportunities for growth. As data continue to play an increasingly important role in business operations, it is likely that more companies will adopt these technologies to stay ahead of the curve and succeed in today's data-driven world [1], [2], [3].

To create a more sustainable and prosperous society by involving advanced technologies, the idea of the future of society, that is, Society 5.0, is proposed. DCAI is a crucial part of this concept because it allows for the collection, analysis, and

Digital Object Identifier 10.1109/MTS.2023.3306532
Date of current version: 21 September 2023.

application of large amounts of data to address complex social and environmental issues. The main objective of DCAI in society is to utilize the potential of big data to create a society that is more efficient, sustainable, and fair as shown in Figure 1. Moreover, DCAI can help to tackle some of the most significant social and environmental challenges that humanity is facing today, such as climate change, poverty, and inequality. By analyzing and modeling complex social and environmental systems, AI can assist policymakers and decision-makers in developing and implementing more effective solutions to these challenges.

In this article, we focus on DCAI and its social impact. For this, several crucial points are highlighted to provide a comprehensive understanding. First, it is essential to understand how DCAI is different from model-centric AI and its potentially transformative effects on various industries, job markets, and society as a whole. We have included case studies and real-world examples of how DCAI can impact communities. Then we emphasize the discussion of the positive social impacts of DCAI, including advancements in healthcare, education, and environmental sustainability. We also give equal importance to understanding the potential challenges as DCAI raises concerns about data privacy, security, and bias, followed by some feasible solution approaches and future directions.

Model-centric versus data-centric AI

Model-centric AI aims to produce an optimal model for a given data set, with the objective of maximizing its performance on a specific task [4]. This is achieved through various model modifications, such as adjusting the loss function, fine-tuning the hyper-parameters, and employing optimization techniques, all aimed at enhancing the model's overall performance, accuracy, precision, and generalization abilities. The accessibility of deep learning models through open-source communities and the availability of paid application programming interfaces (APIs) have made it more convenient for researchers and developers to utilize optimized models across various domains. Despite these advancements, ensuring data quality control remains a persistent challenge. Consequently, placing emphasis on tackling data-related issues becomes crucial to prevent inaccurate outcomes

and enhance the ability of AI models to handle real-world complexities effectively.

In contrast, DCAI focuses on improving the training data set itself to enhance the performance of any given model on a specific AI task. This is accomplished through systematic or algorithmic changes to the data set, ensuring it contains high-quality, relevant, and diverse data during model training [5], [6]. By optimizing the input data, DCAI endeavors to elevate the model's accuracy, robustness, and overall effectiveness in addressing the targeted AI task. DCAI encompasses several techniques, including outlier detection, error correction, data augmentation, feature engineering, establishing consensus labels, active learning, and curriculum learning. These methods aim to enhance data quality, representation, and informativeness to improve AI models' performance and accuracy. The adoption of a DCAI approach extends beyond Big Tech companies, as it proves relevant and beneficial in industries with smaller or more challenging data sets due to regulatory or practical constraints (e.g., manufacturing and healthcare) [7]. For example, OpenAI has openly acknowledged that addressing errors in the training data and labels is a significant challenge with DALL-E and generative pretrained transformer (GPT)-3, underscoring the importance of data quality over the model itself.¹ The improvement of ChatGPT involved fine-tuning the model to prioritize data quality and minimizing harmful, untruthful, or biased output. Moreover, human rankings were incorporated to down-weight "bad data," enhancing the reliability and accuracy of the AI system. Emphasizing data quality plays a pivotal

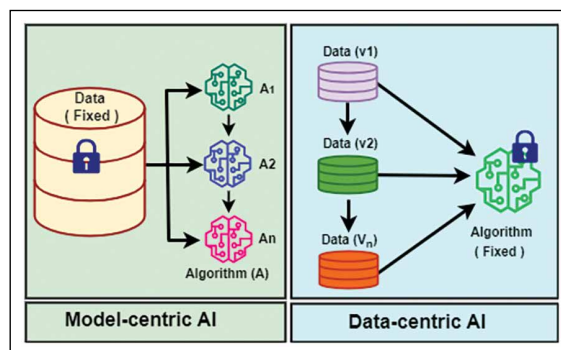


Figure 2. Model-centric and DCAI technology.

¹<https://dcai.csail.mit.edu/>

role in refining AI models and mitigating potential issues related to biased or inaccurate outputs.

The value of model-centric AI and DCAI are not mutually exclusive; instead, they complement each other in building AI systems as shown in Figure 2. Model-centric methods can be used to achieve DCAI goals, such as using generation models like generative adversarial networks (GANs) and diffusion models for data augmentation [6]. Conversely, DCAI can facilitate the improvement of model-centric AI by providing augmented data that inspires advancements in model design. In production scenarios, data and models continuously evolve in an interdependent manner within an environment.

Benefits of data-centric AI

The importance of DCAI for the benefit of society is presented in Figure 1. Maintaining data quality and adapting data to meet evolving needs is essential for maximizing the benefits of DCAI [8]. The industry that adopts DCAI can benefit in the following ways:

More reliable and less biased results

By prioritizing high-quality data, DCAI systems can avoid “garbage-in, garbage-out” and yield more results that are dependable. In addition, the focus on eliminating bias can ensure that the results are fair and equitable.

Lower costs through greater flexibility

DCAI can achieve performance gains with smaller investments in data and resources compared

to model-centric AI. This means organizations can realize cost savings while still improving their AI capabilities.

Less administration

DCAI allows for a more standardized approach that revolves around the data itself, which can reduce the number of models and associated data sets an organization needs to manage. This can lead to reduced costs and greater efficiency [9]. DCAI can help organizations improve their AI capabilities while reducing costs and minimizing bias. By prioritizing high-quality data, organizations can achieve more reliable and accurate results that can lead to better business outcomes [5].

Literature review

Recently, several research articles have proposed DCAI methods in different domains. These include data-centric defense (DCD) for mitigating model inversion attacks [10], crop disease identification in agriculture [11], unsupervised anomaly detection in industrial production [12], transformer-based language models (bidirectional encoder representations from transformers (BERT) and GPT-3 for automated occupation coding [13], enhancing deep neural network (DNN) model robustness [14], a comprehensive community library for biomedical natural language processing (NLP) data sets [15], automatic surgical phase estimation [16], rapid nanoparticle energy prediction [17], identifying incongruous regions in data [18], a multidomain benchmark for

Table 1. Brief summary of related works in DCAI.

Authors	Methodologies	Applications	Limitations
Chen <i>et al.</i> (2023) [10]	Data-Centric Defense	Face recognition	Obscurity of injected samples
Akogo <i>et al.</i> (2023) [11]	YOLO v5X	Crop Disease Identification	Difficulty in distinguishing disease and abiotic classes
Zeiser <i>et al.</i> (2023) [12]	Combination of WGAN and encoder CNN	Automotive component production	No exploration of mixed datasets containing both real and synthetic data
Safikhani <i>et al.</i> (2023) [13]	BERT and GPT3	Social science and official statistics	Limited representation across datasets and languages for job titles and descriptions
Zhong <i>et al.</i> (2022) [14]	DNN, Transferable adversarial	Computer vision	Limited magnitude of improvement, Limited training data, Limited performance gains
Fries <i>et al.</i> (2022) [15]	Multi-task learning	NLP tasks in the biomedical field	Implementation errors may exist in the data loaders, Biases present in the included datasets
Matilla <i>et al.</i> (2022) [16]	ResNet50	Intelligent surgical robots	Encoder overfitting
Kurban <i>et al.</i> (2022) [17]	Random forests, SVM, and KNN	Materials science	Lead to longer processing times
Seedat <i>et al.</i> (2022) [18]	Bayesian and non-Bayesian	Healthcare	Computational complexity, Interpretability, Generalizability
Schmarje <i>et al.</i> (2021) [19]	ImageNet	Image classification	Human annotations in simulations may be incorrect, Large datasets impractical for multiple annotations from an approximated distribution
Baek <i>et al.</i> (2021) [21]	dcSR	Video streaming and content delivery	dcSR requires a server-side component, dcSR may struggle with complex scenes or fast motion

image classification [19], data-centric approaches to improve GAN training with less data [20], and a data-centric super-resolution (dcSR) approach for video quality enhancement [21]. A brief summary of recent DCAI works is presented in Table 1.

Applications of DCAI for the well-being of society

DCAI has the potential to revolutionize many industries and fields by enabling more efficient and effective decision-making based on insights extracted from data. There are countless potential new applications for DCAI, as the field continues to evolve rapidly [22]. A few potential applications for the well-being of society as shown in Figure 3 are briefly discussed below.

Personalized healthcare

Healthcare providers can leverage DCAI to analyze patient data in real time and make more informed treatment decisions. This could involve tailoring drug dosages to specific patients, administering targeted therapies, and predicting diseases with more accuracy. AI algorithms could also examine medical images, recognizing slight variations in tissue or organ structures that might indicate a greater risk of illness. By analyzing genetic data and detecting mutations or genetic markers that may be associated with certain diseases, AI can contribute to identifying effective treatment options. Through the integration of data sets with other sources of information such as medical history and lifestyle factors, AI algorithms can help healthcare providers determine the most appropriate course of action for each patient. This could lead to faster treatment times, more accurate diagnoses, and improved patient outcomes overall [23], [24].

Climate change prediction and mitigation

AI algorithms can be instrumental in analyzing vast quantities of data on weather patterns, greenhouse gas emissions, and other factors to predict future climate change impacts [3]. This can be instrumental in shaping policy decisions and mitigation efforts. With mounting concern over climate change and environmental deterioration, there is a pressing need for precise and timely data on the state of ecosystems and the consequences of human activities. DCAI can scrutinize data from various sources, such as satellite imagery, weather sensors, and wildlife

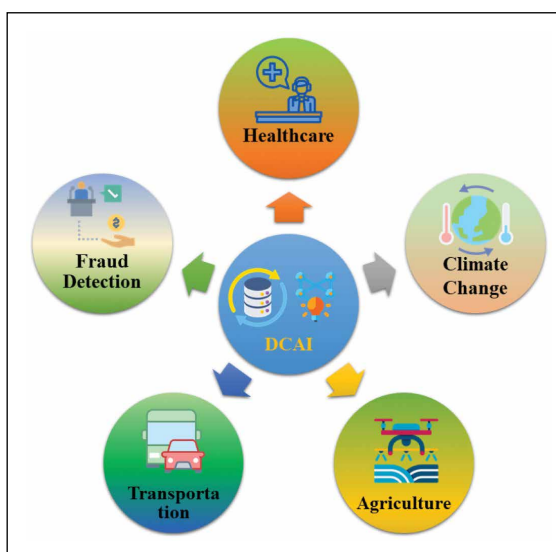


Figure 3. Social applications of DCAI.

monitoring systems, to detect patterns and trends that might signify changes in the environment. This information can be utilized to devise more effective conservation strategies and steer policy decisions.

Intelligent transportation systems

DCAI applied to intelligent transportation systems (ITS) can improve transportation safety, efficiency, and sustainability. By analyzing data from traffic sensors, cameras, global positioning system (GPS) devices, and other sources, AI algorithms can optimize traffic flow, predict and prevent accidents, and manage incidents in real-time [3], [25]. In addition, DCAI can be used to optimize public transportation routes and schedules based on passenger demand and behavior, leading to increased ridership and satisfaction. By combining advanced technologies with DCAI, ITS can become more intelligent and effective, improving the overall transportation experience for users while reducing congestion and environmental impact.

Smart agriculture

DCAI could help farmers optimize crop yields by analyzing data from sensors, weather forecasts, and other sources. This could include real-time decision-making about watering, fertilizing, and harvesting [2]. With the world's population projected to reach 9.7 billion by 2050, food production will need to increase by 70% to meet the growing demand. DCAI can help farmers increase crop yields, reduce waste, and improve sustainability by providing

insights into soil conditions, weather patterns, and crop health. For example, using sensors and drones to collect data on soil moisture levels, temperature, and nutrient content, farmers can leverage DCAI to predict the optimal time to plant, fertilize, and irrigate their crops. They can also use AI-powered predictive models to anticipate weather patterns and adjust their farming practices accordingly, reducing the risk of crop failure due to extreme weather conditions. In addition to improving crop yields and sustainability, DCAI can also help farmers reduce costs by optimizing the use of resources such as water and fertilizer, and by enabling more precise harvesting and processing of crops. The application of DCAI in agriculture has the potential to improve food security, reduce waste, and promote sustainable farming practices, all while increasing profitability for farmers.

Fraud detection and prevention

DCAI is a powerful tool that can be used to prevent fraud across various industries by analyzing large amounts of transaction data to detect patterns of fraudulent activities. AI can identify abnormal behavior, relationships between individuals and organizations, and patterns of fraud using ML algorithms, which can help prevent future fraud. AI can also monitor transactions in real time, assess overall risk, and trigger alerts to prevent fraudulent activity. By applying DCAI to fraud detection and prevention, organizations can detect and prevent fraud more effectively and efficiently, protecting themselves and their customers from the harmful effects of fraud [22].

Opportunities for new ventures and innovation

New ventures and innovations in DC technology have been rapidly growing in recent years. They are leveraging data to develop innovative products and services that are transforming industries and disrupting traditional business models [26], [27]. Some of these are focused on developing advanced analytics solutions to help businesses analyze large data sets and extract insights to make data-driven decisions. AI-based young enterprises are leveraging ML and other advanced techniques to develop intelligent applications that can automate tasks, make predictions, and provide insights. Internet-of-Things (IoT)-oriented startups are developing connected

devices that collect and transmit data, allowing businesses to monitor and optimize their operations. Blockchain-focused startups are developing decentralized applications that use distributed ledgers to store and manage data securely. These applications have the potential to disrupt traditional industries such as finance, supply chain management, and healthcare. Cloud computing is another area where startups are developing infrastructure and platforms that enable businesses to store, process, and analyze data in the cloud [9], [28]. DC technology startups are transforming industries by providing new insights, increasing efficiency, and enabling new business models. As data continue to play an increasingly important role in business and society, these young enterprises are well-positioned to continue driving innovation and growth.

Opportunities in academic research

DCAI education and academic research involve the study of techniques, methods, and algorithms for analyzing, processing, and making predictions based on large data sets. It is a multidisciplinary field that draws upon knowledge from computer science, statistics, mathematics, and engineering [29], [30], [31]. In terms of education, DCAI programs typically cover topics such as data mining, ML, deep learning, NLP, and big data analytics. Students also learn about the various programming languages and tools commonly used in the field, such as Python, R, TensorFlow, PyTorch, and open source tools in DCAI [32], [33], [34] like *Albumentations*—Image augmentation, *Amazon SageMaker Debugger*, *AutoAugment* learning augmentation policies from data, *Bootleg*—self-supervision for named entity disambiguation at the tail, *CleanLab*—clean data set labels, *DataPrep*—prepare your data with a few lines of code, *Deequ*—unit tests for data, *HoloClean*: an ML system for data enrichment, *Knodle*—knowledge-supervised deep learning framework, *Meerkat*—handling data for ML models, *mltrace*—a python package to make ML models observable, *nlpaug*—augment NLP data sets, *Picket*: guarding against corrupted data in tabular data, *Raha & Baran*—configuration free error detection tools, *skweak*—weak supervision for NLP, etc.

Many universities around the world offer degree programs in DCAI, such as bachelor's and master's degrees in data science or AI, as well as PhD programs in related fields. There are also many online courses like MIT Introduction to DCAI and massive

open online courses (MOOCs) available for students who want to learn about DCAI on their own or supplement their formal education. In terms of academic research, DCAI involves developing new algorithms, models, and techniques for analyzing and making predictions based on large data sets [35]. This research can be applied in a variety of fields, such as healthcare, finance, marketing, and social media. Some of the current research topics in DCAI include explainable AI, which seeks to make AI models more transparent and interpretable, and federated learning, which aims to train models in a decentralized manner while maintaining privacy and security. Other research areas include reinforcement learning, transfer learning, etc. [36].

Some challenges of DCAI

Data are often referred to as the new oil, but refining them into meaningful insights can be a challenging process. Many companies are still facing difficulties when it comes to effectively managing and utilizing their data. Despite investing significant time and resources, their efforts often yield little results. The primary challenges are data volume, consistency, quality, and bias that industries face when adopting a DC approach to AI. These challenges can significantly impact the performance and reliability of an AI model, so it is important to address them properly [37], [38]. Organizations must ensure that they have a sufficient volume of high-quality data that is consistent and representative of the data that the model will process after deployment. Blindly collecting as much data as possible can be inefficient and costly. Therefore, it is important to establish the kind of data that is needed before acquiring more data. In addition, organizations must prioritize an effective system of data annotation to ensure consistency in the labeling of the data. An AI model trained on inconsistent data annotation can quickly become unreliable.

Biases can be introduced into AI systems when the training data used to develop the models are themselves biased. If the training data is not diverse or representative enough, or if it reflects historical biases and stereotypes, then the resulting AI models may perpetuate or even amplify these biases. In addition, biases can also be introduced during the selection of features or the design of the algorithm used to develop the AI model. Therefore, it is important to carefully select and prepare training data and

algorithmic models to minimize the potential for biases in AI systems [39], [40]. For example, in 2018, Amazon had to abandon its AI-based hiring tool, which was designed to help automate and streamline the recruitment process. The tool was trained on resumes submitted to the company over a 10-year period, which were predominantly from men as the tech industry is male-dominated. As a result, the tool learned to prefer male candidates over female candidates, reflecting the biases present in the training data. Amazon's tool was ultimately not used because it perpetuated the existing gender bias in the tech industry, rather than mitigating it. This case highlights the importance of ensuring that training data is diverse and representative and that AI models are tested for fairness to avoid perpetuating or amplifying biases [41], [42].

Few feasible solution approaches

The solution to DCAI challenges requires a multifaceted approach. Data sources should be diversified to ensure the algorithm is exposed to a more representative sample of the population. Algorithms should be audited regularly to identify biases and provide opportunities for correction. Ethical considerations should be included in the development process [3]. Human oversight can help identify potential biases and ensure ethical decision-making. Monitoring and evaluating the algorithm's performance can help identify and mitigate biases over time, ensuring reliable and ethical results.

Transparency

Transparency in AI decision-making processes is crucial for several reasons: When AI systems make decisions that affect people's lives, it's important to know how those decisions are being made. Transparency allows for the identification of biases or errors in the decision-making process, and it helps ensure that those responsible for the AI systems can be held accountable for their actions [3]. For building trust between users and AI systems, transparency plays a key role. If users don't know how an AI system is making decisions, they are unlikely to trust the system. On the other hand, if the decision-making process is transparent, users are more likely to trust the system, which can lead to greater adoption and use.

Transparency can help to ensure that AI systems are making decisions fairly. By making the decision-making process transparent, it is easier

to identify and address biases that may be present in the system. This can help ensure that the system is making decisions that are fair and unbiased. It's also important to be able to explain how a decision was made by an AI system. This can help users understand the reasoning behind the decision and can provide important context for how the system works [43].

Therefore, transparency is critical for ensuring that AI systems are making decisions that are fair, unbiased, and trustworthy. It can also help build user trust and provide important context for how the system works. To increase transparency in DCAI systems, it is important to open-source the code and data, provide clear explanations for model outputs, and disclose training data sources. This helps identify potential biases and errors, build user trust, and provide context for the system's decision-making process.

Data diversity and representative

The use of diverse and representative data sets is important to reduce biases in AI systems. When AI systems are trained on biased or unrepresentative data, they may make biased decisions, which can have negative impacts on marginalized groups [44]. By using diverse and representative data sets, AI systems can be trained to make more fair and unbiased decisions that better reflect the real world. There are several techniques for improving data diversity and representativeness, including:

Data augmentation

This involves creating additional training data by applying transformations to the existing data set. For example, flipping images horizontally or adding noise to audio data. This can help increase the diversity of the data set and make it more representative of the real world [45].

Active learning

This involves selecting the most informative samples for human labeling, which can help improve the representativeness of the data set. By selecting samples that are most likely to reduce model uncertainty or increase diversity, active learning can help ensure that the labeled data set is more representative of the real world [46].

Crowd-sourcing

This involves outsourcing the labeling of data to a large group of people, which can help improve the diversity and representativeness of the data set. Crowdsourcing can help ensure that the labeled data set reflects a wider range of perspectives and experiences, which can help reduce biases.

Ongoing monitoring and evaluation of data sets are important to ensure continued diversity and representativeness. Over time, the real world may change, and new types of data may become more relevant [45], [47]. If data sets are not regularly monitored and evaluated, they may become less diverse and less representative of the real world, which can lead to biased AI systems. By regularly monitoring and evaluating data sets, AI systems can be trained on the most current and representative data, which can help reduce biases and ensure fair and equitable decision-making [48].

Fairness and bias testing

AI models need to be tested for fairness and biasedness to ensure that they are not making decisions that unfairly advantage or disadvantage certain groups of people. Without testing, it is possible for AI models to perpetuate social biases and inequalities [43]. By testing for fairness and biasedness, AI models can be improved and made more equitable, ensuring that they serve all users fairly and impartially.

Selecting appropriate fairness metrics and performing comprehensive bias tests are important to ensure that AI models are evaluated fairly and accurately. Fairness metrics should align with the specific context of the AI model and the potential impact on different groups of people. Bias tests should be performed to identify potential biases in the data and ensure that the model does not unfairly discriminate against certain groups. By selecting appropriate fairness metrics and performing comprehensive bias tests, AI models can be evaluated in a thorough and fair manner, helping to reduce biases and promote equity. There are several techniques for testing data-centric approaches, including:

Sensitivity analysis

This involves testing how sensitive the model is to changes in the input data, which can help identify potential biases and sources of error.

Group fairness testing

This involves testing whether the model is treating different groups of people fairly and equally, and can help identify potential biases that may be present in the model.

Counterfactual analysis

This involves testing how the model's decisions would change if certain variables were altered, which can help identify potential sources of bias and help improve the model's fairness.

These techniques can help identify potential biases and sources of error in AI models, which is important for ensuring that the models are fair and equitable [3], [43].

Future directions

The future of DCAI technology is exciting and holds enormous potential. With advances in AI algorithms, ML models, and data analytics tools, companies can gain deeper insights into their operations, customers, and markets. This allows them to make more informed decisions, optimize processes, and create more personalized and engaging customer experiences. One of the most significant trends in DCAI technology is the increased use of NLP and conversational AI. NLP enables machines to understand human language and respond accordingly, opening up new possibilities for customer service, virtual assistants, and chatbots. Conversational AI takes this a step further by allowing machines to carry on a conversation with a person in a more natural and human-like way [30], [33]. Another trend in DCAI technology is the integration of AI and IoT devices. This allows companies to collect and analyze data from a wide range of sources, including sensors, machines, and other devices. These data can then be used to optimize operations, improve product design, and create more personalized experiences for customers [34], [37]. Finally, there is the potential for DCAI technology to make significant strides in healthcare, financial services, and other industries. AI can help healthcare providers diagnose diseases more accurately, develop new treatments, and improve patient outcomes. In the financial sector, AI can help identify fraud, streamline lending processes, and improve risk management.

AS TECHNOLOGY AND DCAI continue to evolve, it holds the promise of transforming our lives, work experiences, and interactions in unprecedented

ways. Its current applications in healthcare, education, and finance demonstrate its ability to improve outcomes and enhance human experiences. However, despite the immense potential, there are significant concerns regarding the impact of DCAI, such as privacy infringement, job displacement, and biased decision-making. Therefore, it is crucial to develop and deploy DCAI based on ethical principles to ensure the equitable distribution of benefits across society. ■

References

- [1] J. He et al., "The practical implementation of artificial intelligence technologies in medicine," *Nature Med.*, vol. 25, no. 1, pp. 30–36, Jan. 2019.
- [2] C. Zhang and Y. Lu, "Study on artificial intelligence: The state of the art and future prospects," *J. Ind. Inf. Integr.*, vol. 23, Sep. 2021, Art. no. 100224.
- [3] A. Rawal et al., "Recent advances in trustworthy explainable artificial intelligence: Status, challenges, and perspectives," *IEEE Trans. Artif. Intell.*, vol. 3, no. 6, pp. 852–866, Dec. 2022.
- [4] O. H. Hamid, "From model-centric to data-centric AI: A paradigm shift or rather a complementary approach?" in *Proc. 8th Int. Conf. ITT*, May 2022, pp. 196–199.
- [5] A. Ng, "AI does not have to be too complicated or expensive for your business," *Harvard Bus. Rev.*, Jul. 2021. [Online]. Available: <https://hbr.org/2021/07/ai-doesnt-have-to-be-too-complicated-or-expensive-for-your-business>
- [6] D. Zha et al., "Data-centric artificial intelligence: A survey," 2023, *arXiv:2303.10158*.
- [7] M. H. Jarrahi, A. Memariani, and S. Guha, "The principles of data-centric AI (DCAI)," 2022, *arXiv:2211.14611*.
- [8] S. E. Whang et al., "Data collection and quality challenges in deep learning: A data-centric AI perspective," *VLDB J.*, vol. 32, no. 4, pp. 791–813, Jul. 2023.
- [9] X. Zhang et al., "SliceTeller: A data slice-driven approach for machine learning model validation," *IEEE Trans. Vis. Comput. Graphics*, vol. 29, no. 1, pp. 842–852, Jan. 2023.
- [10] C. E. Si, "Data-centric defense: Shaping loss landscape with augmentations to counter model inversion," in *Proc. 40th ICML*, 2023, pp. 1–3.
- [11] A. Darlington et al., "Localized data work as a precondition for data-centric ML: A case study of full lifecycle crop disease identification in Ghana," in *Proc. 40th ICML*, 2023, pp. 1–14.

- [12] A. Zeiser et al., "Evaluation of deep unsupervised anomaly detection methods with a data-centric approach for on-line inspection," *Comput. Ind.*, vol. 146, Apr. 2023, Art. no. 103852.
- [13] P. Safikhani et al., "Automated occupation coding with hierarchical features: A data-centric approach to classification with pre-trained language models," *Discover Artif. Intell.*, vol. 3, no. 1, p. 6, Feb. 2023.
- [14] Y. Zhong et al., "Exploiting the potential of datasets: A data-centric approach for model robustness," in *Proc. AAAI Workshop Adversarial Mach. Learn.*, 2022, pp. 1–5.
- [15] J. Fries et al., "BigBio: A framework for data-centric biomedical natural language processing," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 25792–25806.
- [16] R. Sanchez-Matilla et al., "Data-centric multi-task surgical phase estimation with sparse scene segmentation," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 17, no. 5, pp. 953–960, May 2022.
- [17] H. Kurban, M. Kurban, and M. M. Dalkilic, "Rapidly predicting Kohn–Sham total energy using data-centric AI," *Sci. Rep.*, vol. 12, no. 1, p. 14403, Aug. 2022.
- [18] N. Seedat, J. Crabbé, and M. van der Schaar, "Data-SUITE: Data-centric identification of in-distribution incongruous examples," in *Proc. 39th ICML*, 2022, pp. 19467–19496.
- [19] L. Schmarje, Y.-H. Liao, and R. Koch, "A data-centric image classification benchmark," in *Proc. NeurIPS Datacentric AI Workshop*, 2021, p. 7.
- [20] M. Motamedi, N. Sakharnykh, and T. Kaldewey, "A data-centric approach for training deep neural networks with less data," in *Proc. 35th Conf. NeurIPS*, 2021, pp. 1–6.
- [21] D. Baek et al., "DcSR: Practical video quality enhancement using data-centric super resolution," in *Proc. 17th Int. Conf. Emerg. Netw. Exp. Technol.*, Dec. 2021, pp. 336–343.
- [22] S. Bergamaschi et al., "Big data integration for data-centric AI," in *Proc. 1st ItaData*, 2022, p. 1.
- [23] L. Oberste and A. Heinzl, "User-centric explainability in healthcare: A knowledge-level perspective of informed machine learning," *IEEE Trans. Artif. Intell.*, vol. 4, no. 4, pp. 840–857, Aug. 2023.
- [24] J. Waring, C. Lindvall, and R. Umerton, "Automated machine learning: Review of the state-of-the-art and opportunities for healthcare," *Artif. Intell. Med.*, vol. 104, Apr. 2020, Art. no. 101822.
- [25] V. Lakshmanan, *Data Science on the Google Cloud Platform*. Sebastopol, CA, USA: O'Reilly Media, 2022.
- [26] R. C. Fernandez et al., "Aurum: A data discovery system," in *Proc. IEEE 34th ICDE*, Apr. 2018, pp. 1001–1012.
- [27] M. Cafarella et al., "Ten years of webtables," *Proc. VLDB Endowment*, vol. 11, no. 12, pp. 2140–2149, Aug. 2018.
- [28] D. Brickley, M. Burgess, and N. Noy, "Google dataset search: Building a search engine for datasets in an open web ecosystem," in *Proc. World Wide Web Conf.*, May 2019, pp. 1365–1375.
- [29] V. S. Sheng, F. Provost, and P. G. Ipeirotis, "Get another label? Improving data quality and data mining using multiple, noisy labelers," in *Proc. 14th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2008, pp. 614–622.
- [30] M. Dreves et al., "Validating data and models in continuous ML pipelines," *IEEE Data Eng. Bull.*, vol. 44, no. 1, pp. 42–50, Mar. 2021.
- [31] E. Breck et al., "Data validation for machine learning," in *Proc. Mach. Learn. Syst.*, vol. 1, 2019, pp. 334–347.
- [32] T. Rekatsinas et al., "HoloClean: Holistic data repairs with probabilistic inference," *Proc. VLDB Endowment*, vol. 10, no. 11, pp. 1190–1201, Aug. 2017.
- [33] P. Li et al., "CleanML: A benchmark for joint data cleaning and machine learning [experiments and analysis]," in *Proc. 37th IEEE Int. Conf. Data Eng.*, Jun. 2021, pp. 13–24.
- [34] S. Krishnan et al., "ActiveClean: Interactive data cleaning for statistical modeling," *Proc. VLDB Endowment*, vol. 9, no. 12, pp. 948–959, Aug. 2016.
- [35] L. Jiang et al., "MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2304–2313.
- [36] B. Han et al., "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1–11.
- [37] P. W. Koh, J. Steinhardt, and P. Liang, "Stronger data poisoning attacks break data sanitization defenses," *Mach. Learn.*, vol. 111, no. 1, pp. 1–47, Jan. 2022.
- [38] A. Paudice et al., "Detection of adversarial training examples in poisoning attacks through anomaly detection," *CoRR*, pp. 1–10, 2018.
- [39] M. Feldman et al., "Certifying and removing disparate impact," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2015, pp. 259–268.
- [40] M. B. Zafar et al., "Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment," in *Proc. 26th Int. Conf. World Wide Web*, Apr. 2017, pp. 1171–1180.

- [41] A. Chouldechova, "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments," *Big Data*, vol. 5, no. 2, pp. 153–163, Jun. 2017.
- [42] C. Dwork et al., "Fairness through awareness," in *Proc. 3rd Innov. Theor. Comput. Sci. Conf.*, Jan. 2012, pp. 214–226.
- [43] W. Liang et al., "Advances, challenges and opportunities in creating data for trustworthy AI," *Nature Mach. Intell.*, vol. 4, no. 8, pp. 669–677, Aug. 2022.
- [44] D. Berthelot et al., "ReMixMatch: Semi-supervised learning with distribution matching and augmentation anchoring," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–13.
- [45] E. D. Cubuk et al., "AutoAugment: Learning augmentation strategies from data," in *Proc. IEEE/CVF Conf. CVPR*, Jun. 2019, pp. 113–123.
- [46] B. Settles, "Active learning literature survey" Dept. Comput. Sci., Univ. Wisconsin–Madison, Madison, WI, USA, Tech. Rep. 1648, Jan. 2010.
- [47] D. Berthelot et al., "MixMatch: A holistic approach to semi-supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–11.
- [48] D. Hendrycks et al., "AugMix: A simple data processing method to improve robustness and uncertainty," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–15.

Sushant Kumar is a research associate III at the Department of Computer Science and Engineering, Indian Institute of Technology (IIT BHU), Varanasi 221005, India. His research interests include joint compressed sensing, dictionary learning, and ML for biomedical signal analysis.

Ritesh Sharma is a PhD scholar at the Department of Computer Science and Engineering, Indian Institute of Technology (IIT BHU), Varanasi 221005, India. He has published several peer-reviewed journal articles. His research interests are deep learning, ML, and bioinformatics.

Vishakha Singh is a prime minister research fellow pursuing a PhD at the Department of Computer Science and Engineering, Indian Institute of Technology (IIT BHU), Varanasi 221005, India. She has authored and coauthored several articles in peer-

reviewed high-impact factor journals. Her research interests include artificial intelligence, deep learning, machine learning, and multi-objective optimization.

Shrikant Tiwari is an associate professor at the Department of Computer Science and Engineering (CSE), Institute of Engineering and Rural Technology (IERT), Prayagraj 211002, India. His current research interests include machine learning, deep learning, computer vision, medical image analysis, pattern recognition, and biometrics. He is a Member of ACM, IET, FIETE, CSI, ISTE, IAENG, and SCIEI, reviewer for many international journals of repute, and a Senior Member of IEEE.

Sanjay Kumar Singh is a professor at the Department of Computer Science and Engineering, Indian Institute of Technology (IIT BHU), Varanasi 221005, India. He has authored or coauthored more than 150 national and international journal publications, book chapters, and conference papers. His current research interests include machine learning, deep learning, computer vision, medical image analysis, pattern recognition, and biometrics. He is also a guest editorial board member and a reviewer for many international journals of repute. He is a Senior Member of IEEE.

Sumit Datta is an assistant professor at the School of Electronic Systems and Automation, IIITM-Kerala/Digital University Kerala, Thiruvananthapuram 695317, India. His research interests include biomedical signal/image processing, compressed sensing MRI, super-resolution, and medical image analysis using deep learning. He is also a member of editorial board and a reviewer for many international journals.

■ Direct questions and comments about this article to Sumit Datta, School of Electronic Systems and Automation, IIITM-Kerala/Digital University Kerala, Thiruvananthapuram 695317, India; sumit.datta@iiitm.ac.in.