Article

Patterns

Unveiling value patterns via deep reinforcement learning in heterogeneous data analytics

Graphical abstract

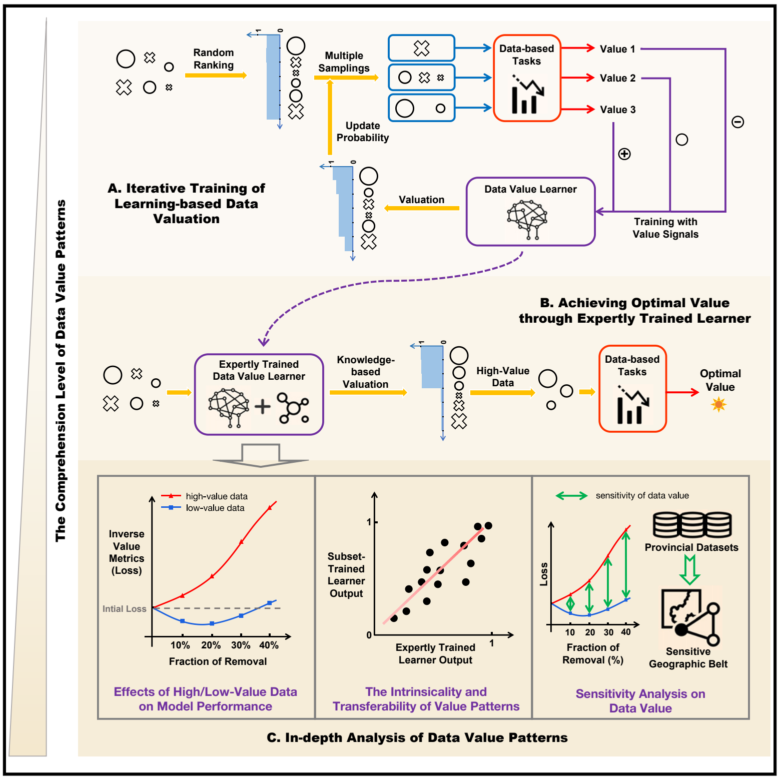

Highlights

- We introduce a learning-based paradigm for data valuation across scenarios
- Our study identifies varying effects of high/low-quality data on model efficacy
- This method explores inherent and transferable value patterns across datasets
- Analysis reveals data value's geographic sensitivity in nationwide power forecasts

Authors

Yanzhi Wang, Jianxiao Wang, Feng Gao, Jie Song

Correspondence

wang-jx@pku.edu.cn (J.W.),
jie.song@pku.edu.cn (J.S.)

In brief

This study presents a learning-based method for evaluating data quality tailored to AI applications, underlining the critical role of data value in boosting big-data performance. This approach can enhance model efficiency, data usage, and governance. This work seeks to reconcile data quality with its practical applications, offering a streamlined path from data valuation to application in data-driven fields.

Wang et al., 2024, Patterns 5, 100965
May 10, 2024 © 2024 The Author(s).
https://doi.org/10.1016/j.patter.2024.100965

Patterns

CellPress
OPEN ACCESS

Article

Unveiling value patterns via deep reinforcement learning in heterogeneous data analytics

Yanzhi Wang,[1] Jianxiao Wang,[2,3,4,*] Feng Gao,[1] and Jie Song[1,2,3,*]
[1]Department of Industrial Engineering and Management, College of Engineering, Peking University, Beijing 100871, China
[2]National Engineering Laboratory for Big Data Analysis and Applications, Peking University, Beijing 100871, China
[3]PKU-Changsha Institute for Computing and Digital Economy, Changsha 410000, China
[4]Lead contact
*Correspondence: wang-jx@pku.edu.cn (J.W.), jie.song@pku.edu.cn (J.S.)
https://doi.org/10.1016/j.patter.2024.100965

THE BIGGER PICTURE In the era of big data, the surge in volume is matched by the challenges of data quality, which resonates across all domains of data-driven and artificial-intelligence-related technologies. This research proposes a method to navigate these challenges by introducing a paradigm based on deep reinforcement learning, capable of discerning value in data across varied contexts. By enabling the strategic selection of optimal data samples, we envision a future where analytics are not just smarter but also more adaptable, allowing decision makers to harness the full potential of their data assets. The implications of this work extend beyond technical realms, offering insights that could shape policies and provide fresh perspectives in data-driven industries.

SUMMARY

Artificial intelligence has substantially improved the efficiency of data utilization across various sectors. However, the insufficient filtering of low-quality data poses challenges to uncertainty management, threatening system stability. In this study, we introduce a data-valuation approach employing deep reinforcement learning to elucidate the value patterns in data-driven tasks. By strategically optimizing with iterative sampling and feedback, our method is effective in diverse scenarios and consistently outperforms the classic methods in both accuracy and efficiency. In China's wind-power prediction, excluding 25% of the overall dataset deemed low-value led to a 10.5% improvement in accuracy. Utilizing just 42.8% of the dataset, the model discerned 80% of linear patterns, showcasing the data's intrinsic and transferable value. A nationwide analysis identified a data-value-sensitive geographic belt across 10 provinces, leading to robust policy recommendations informed by variances in power outputs and data values, as well as geographic climate factors.

INTRODUCTION

Artificial intelligence (AI) technologies have revolutionized the utilization of data for constructing machine-learning models, thereby empowering and optimizing real-world research analysis and production processes, leading to tangible benefits.[1–4] For example, in the low-carbon energy revolution, the integration of digital technologies links the production, transmission, consumption, and storage of renewable energy with data and control systems, building the basic support of the modern energy industry.[5–7] However, the sheer volume and diversity of data not only lead to an abundance of low-quality data with escalating processing and storage costs but also pose challenges in training efficiency due to the variable significance of data in AI-based modeling. The existence of low-quality data causes models to learn complex and irrelevant numerical rules. This causes a decrease in the accuracy and the efficiency of these methods in the absence of data screening.[8–11] The instability in data quality resulting from the drastic and unpredictable nature of natural climatic conditions poses challenges in comprehending renewable-energy patterns. This instability jeopardizes the safety and the reliability of power systems.[12–14] By discerning the varying contributions of data of different qualities to practical scenarios, informed decisions can be made regarding the retention or elimination of data based on their quality. This strategic approach to extracting and utilizing "smart data"—the data that are most relevant and beneficial for specific tasks—enhances the refinement of data-driven models while simultaneously maximizing the overall value of the data.[15,16] The criteria for data screening vary significantly depending on the specific application of the data. This variability highlights the urgent need for a flexible, comprehensive data-valuation paradigm.

Such a paradigm is essential for advancing data-driven modeling across diverse scenarios.

However, the research on data-value assessment is currently limited. Most existing approaches for evaluating dataset quality lack unified quantitative measurements and supervision methods that incorporate actual scenario model outputs.[17,18] Moreover, these approaches do not reliably benchmark the practical usage of data. While some statistical filtering algorithms effectively denoise datasets by identifying outliers, thus improving the distribution for enhanced model generalization,[19] these binary classifications fall short in differentiating the nuances of various data points for precise dataset adjustments. More crucially, such methods overlook the influence of usage scenarios on data-value judgment, such as the contribution of outliers to model robustness in extreme-event models. Scholars use the Shannon entropy and the non-noise ratio to describe data quality[20] and to explain the value of data through the reduction in uncertainty[21] on renewable energy. However, neither metric is used to explain value creation at the data level. From this standpoint, while the leave-one-out (LOO) method effectively evaluates the marginal utility of an individual data point against the whole dataset, its ability to extend this evaluation to the cumulative value of data is limited, which restricts its use in broader big-data analysis. Considering the complex interrelations among data points, some scholars liken them to players in cooperative games by applying game-theoretical metrics such as the Shapley value.[22] However, despite machine-learning enhancements for the algorithm,[23] the applicability of the Shapley value to large-scale data and complex models remains limited. Furthermore, these methods often exhibit biases in selecting data subsets that maximize value, a limitation stemming from the averaging calculations they rely on. Some scholars have advanced the algorithms by integrating meta-learning and deep-learning principles,[24] primarily in homogeneous data structures such as label diagnosis and image recognition, where valuation involves discrete, regular samples.[25] However, in fields such as renewable energy, which are characterized by time-series data, the continuous, intermittent, and complex nature of these datasets necessitates a valuation approach that is simultaneously flexible and specialized.[26] A general valuation approach should be adaptable and robust to address the potential uncertainties and the instabilities in practical data usage and should be effectively tailored to meet the unique challenges of each specific domain. Furthermore, the role of data valuation extends beyond conventional applications, especially in scenarios with varied objectives. For instance, in real-world analytics, data are crucial for guiding optimal decision making. This drives the need for algorithms that are compatible with a broader spectrum of value representations that go beyond simple predictions.[27,28] This necessity highlights the importance of developing a comprehensive and efficient framework for data valuation that can address the multifaceted and intricate nature of data-value assessment across diverse contexts.

In this research, we establish an effective data-valuation paradigm suitable for a broad spectrum of data-driven scenarios. Within this framework, the values derived from continuously sampled data subsets are interpreted as feedback signals, progressively revealing the intrinsic patterns that determine the value of the data. By applying our methodology to various data-driven contexts—including tasks such as adult income classification, forest fire and obesity regression, and heart failure clustering—we validate the efficacy of our algorithm and demonstrate its adaptability across various data types, models, and performance metrics. These metrics include the accuracy, mean absolute error (MAE), mean squared error (MSE), and sum of distances to cluster centers (distance), which gauges dataset dispersion. Our numerical results indicate that this learning-based approach significantly outperforms traditional methods such as the LOO method and the Shapley value in terms of both data-value calculation efficacy and operational speed. We further apply our framework to the analysis of wind-power data in China, highlighting its practical utility in renewable-energy prediction scenarios. This application not only demonstrates the effectiveness of the method but also aids in uncovering data patterns that can inform policy recommendations. The findings of this study are particularly noteworthy for the exceptional precision in handling uncertainty across diverse datasets. The results of this study demonstrate the inherent and transferable nature of data-value patterns, with a special focus on the geographic variations in data-value sensitivity across Chinese provinces. These insights, supported by an integration of power-based and geographic knowledge, provide a foundation for informed, data-driven regulation recommendations.

RESULT

Learning-based data valuation

In this section, we explore a learning-based approach to data valuation by using a deep neural network. The network functions as a value learner, processing individual data features to estimate their respective values. This method utilizes the ability of deep learning to discern intricate patterns within the dataset, facilitating precise data-value estimations. The primary goal is to identify a data subset that optimizes the utility, a measure, derived from the dataset, of the performance of the data-based task. The value of each data point is quantified by its probability of inclusion in this optimal subset, with values near 1 indicating a greater likelihood of selection due to their substantial impact on utility enhancement.

Figure 1 illustrates the iterative training of the value learner, beginning with an equalized value distribution across the data, represented as a random ranking. This iterative process involves repeatedly sampling data subsets based on their current value estimations within each cycle. This method not only augments the stability of the learning process but also optimizes the efficiency of sample utilization. During each iteration, the utility of the selected subsets is assessed in the context of a specific data-driven task. This assessment yields either positive or negative feedback. This feedback is used to adjust the data-valuation strategy, ultimately guiding the learner to prioritize data subsets with maximal utility potential and to devalue fewer contributive points. In summary, the training process enables the learner to gradually master the strategy of selecting the subset with the highest utility, thereby enhancing model performance and the decision-making capabilities in terms of data valuation.

In practical scenarios, especially in complex fields such as wind-power prediction, data-value expression often transcends binary categorizations (merely 1s and 0s) due to

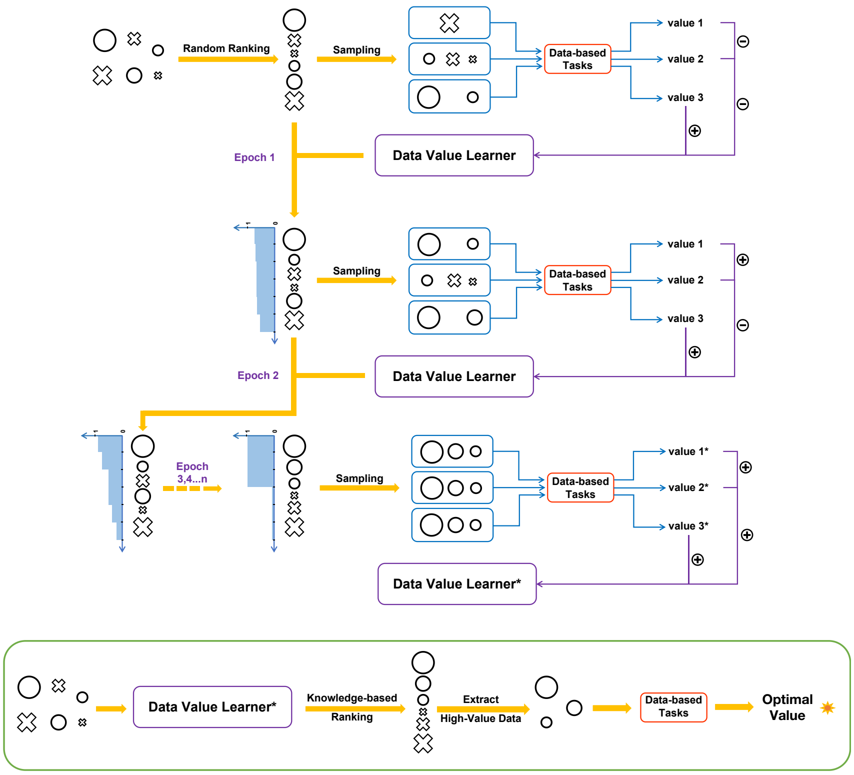

Figure 1. LDV paradigm
This figure illustrates our approach, where we use the term "data-based tasks" to emphasize the incorporation of various valuation metrics applicable in data-driven contexts. For simplicity in this visualization, we assume that the data elements contribute independently to the task. Here, circles and crosses symbolize positive and negative contributions, respectively. Plus and minus signs indicate the impact of sampled subsets on the value performance of the task, either enhancing or diminishing the learner's perceived value for these combinations. The blue bars graphically represent the data-value outputs assigned by the trained learner to each data point in their respective positions. An asterisk symbolizes the optimal learner state achieved post training. The green box depicts the final application phase, where the optimized learner is operationalized. Let us note, however, that real-world scenarios may present more intricate interdependencies among data, leading to a more complex distribution of data values than is depicted here.

intricate interrelations such as substitution and complementarity effects among data points. Consequently, data values can manifest as continuous variables that represent the probability of their inclusion in the optimal subset rather than as discrete entities. This continuous nature reflects a nuanced understanding of data value, allowing the data-value learner to generate a distinctive value distribution. This perspective accurately captures the subtleties and complexities inherent in real-world data relationships.

The cornerstone of this learning-based approach is utility feedback obtained from the data subsets. Its general applicability across various data-driven contexts is attributed to the focus on utility estimation. This data-valuation approach necessitates mapping only between data subsets and their utility values, regardless of task complexity or utility metrics. However, a challenge arises due to the absence of gradient information when discretizing the continuous value output of the data learner for the subset in the sampling mechanism. This issue blocks the training process of the neural network, particularly during backpropagation. To overcome this hurdle and to ensure theoretical robustness, the model employs a policy gradient algorithm within a deep reinforcement-learning framework.[29] This strategy redefines the learner's objective as an optimization within the context of reinforcement learning. The algorithm plays a crucial role in making the neural network trainable while expediting the adaptation and refinement of the data-value strategy. To further enhance the training process, techniques such as importance sampling and clipping functions[30–33] are implemented. These methods are instrumental in promoting algorithm convergence and stability, thereby significantly improving the reliability and the effectiveness of the data-valuation model.

Experiments for various data-driven tasks

To evaluate the adaptability and the effectiveness of our learning-based data-valuation (LDV) algorithm, we conducted experiments across four distinct datasets from various disciplines, namely the social, natural, life, and medical sciences. These datasets, which are critical to a range of practical tasks, include adult income level classification by using census features,[34] forest fire size prediction with meteorological and geographic data,[35] obesity-level estimation based on dietary habits and physical conditions,[36] and clustering analysis of clinical features in heart-failure patients.[37] For these diverse tasks, we selected appropriate machine-learning methods and assessment metrics, as detailed in Table 1. Our comparative analysis focused primarily on the LOO and Shapley value (SV) methods; we evaluated their performance and computational efficiency.

The precision of the data valuation for each method was assessed by sorting and excluding a predefined fraction of the data based on their calculated values. This was followed by reassessing the utility of the remaining subset. Specifically, we removed the data in both ascending and descending order of value ranking and re-evaluated the model performance with the condensed dataset. Optimal valuation significantly enhances the final value by accurately identifying and removing low-value data, while the elimination of high-value points leads to a marked decrease in value. This indicates the precise representation of essential data features.

In Table 1, biases in data-value calculations when using LOO and SV are identified; these biases affect the performance in certain tasks. For instance, SV fails to reduce the MSE in the obesity dataset, while LOO fails to increase accuracy in census classification. In contrast, LDV consistently yields significant improvements across datasets, tasks, and valuation methods, outperforming LOO and SV. These improvements are substantial, with LDV performing 6.76 times better than LOO does in the

Table 1. Performance comparison under various datasets, models, and value metrics

Subject	Task Model Metric	Method	Promotion (within 40% removal)	Declination (within 40% removal)	Running time (s)
Adult census	classification	LOO	–	−8.5%	22
	XGB	SV	+0.5%	−11.5%	126
	accuracy	LDV	+1.1%*[1]	−15.4%*	104
Forest fire	regression	LOO	+0.025	−0.056	5
	SVR	SV	+0.089	−0.053	562
	MAE reduction	LDV	+0.169*	−0.105*	103
Obesity	regression	LOO	+0.021	−0.388	102
	LGBM	SV	–	−0.487	2,448
	MSE reduction	LDV	+0.031*	−0.611*	387
Heart failure	cluster	LOO	+0.068	–	2
	K-means	SV	+0.070	−0.009	320
	distance reduction	LDV	+0.264*	−0.013*	52

[1]Note: An asterisk (*) indicates the method that achieved the best performance.

forest fire dataset and 3.77 times better than SV does in the heart-failure dataset. This remarkable improvement is also evident in the decrease in the metrics upon removal of high-value data, where the LDV surpasses the LOO and SV algorithms in terms of performance. For instance, in the census dataset, the LDV outperforms the LOO by 6.9% and the SV by 3.9%.

Regarding computational efficiency, LOO has a shorter execution time due to its deterministic and straightforward calculation process, albeit at the cost of valuation precision. Despite using Monte Carlo methods to expedite SV computations, achieving accurate data-valuation convergence remains time intensive. Conversely, LDV rapidly achieves optimal data valuations through a reinforcement-learning framework, effectively combining exploration and greedy strategies in agent network training. As tasks and metrics grow more complex, the time advantage of the LDV becomes more pronounced. For instance, in regression problems, the SV execution times are 5.46 and 6.33 times longer, respectively. In clustering tasks with flexible utility definitions, LDV requires only 0.16 times the time taken by SV. Although LDV has a longer runtime than LOO does, its computational cost is reasonable; this is especially true when considering its superior accuracy in data valuation compared to that of the former.

In Figure 2, metric inversion is applied to the vertical axis to effectively align the trend curves. For the LDV classification (Figure 2A) and regression tasks (Figures 2B and 2C), we observed a notable rebound effect. Initially, removing low-value data containing redundant information enhances model performance; this enhancement demonstrates efficient data governance. However, as progressively more data are eliminated and the value of the removed data increases, crucial feature data starts being excluded. This results in a subsequent performance decline, underscoring the necessity of maintaining a diverse dataset for unbiased training and optimal utility in test environments.

In clustering tasks (Figure 2D), the pattern deviates due to the unsupervised nature of these tasks. Clustering lacks the stringent requirement for specific feature data essential for performance validation in supervised tasks. Therefore, no comparable rebound phenomenon is observed. The LDV algorithm effectively identifies and removes isolated data points. This process accelerates the formation of clustering centers. In contrast, LOO and SV are slower at recognizing data contributions in clustering tasks, resulting in a more gradual formation of cluster centers. This distinction underscores the adaptability and the effectiveness of LDV in addressing the distinct requirements of both supervised and unsupervised learning environments. Furthermore, these aspects of the model help to elucidate the governing principles of data value from two distinct perspectives, enhancing our understanding of the intricate role of data across various machine-learning contexts.

In Figures 2A, 2B, and 2C, the removal of high-value data in the LDV exhibits a consistent pattern: as more data are excluded, the inverse value metric correspondingly increases. Over the long term, LDV shows a faster and more stable rate of increase in inverse value than do other methods, highlighting the direct impact of removing high-value data on the utility and efficacy of the model training for the dataset. The removal process, starting with the most critical features, leads to a direct and steady decline in performance. This downward trajectory persists even as the value of the removed data decreases, indicating that the training of the model deteriorates continuously with decreasing data, regardless of the diminishing value of the data being excluded.

In Figure 2D, the ability of LDV to maintain the discreteness of data points notably surpasses that of the other methods, illustrating the effectiveness of our data-valuation framework. Overall, beyond the exemplary performance of LDVs in data-removal experiments, the data-value patterns theoretically align with the principles of machine learning. This alignment is significant because it demonstrates the robust understanding and effectiveness of LDV in terms of data valuation. In contrast to the unstable data-value expression observed in LOO and SV, the LDV approach more accurately reflects our expectations of data-value patterns, affirming its superior performance in both comprehension and practical application within data-driven models.

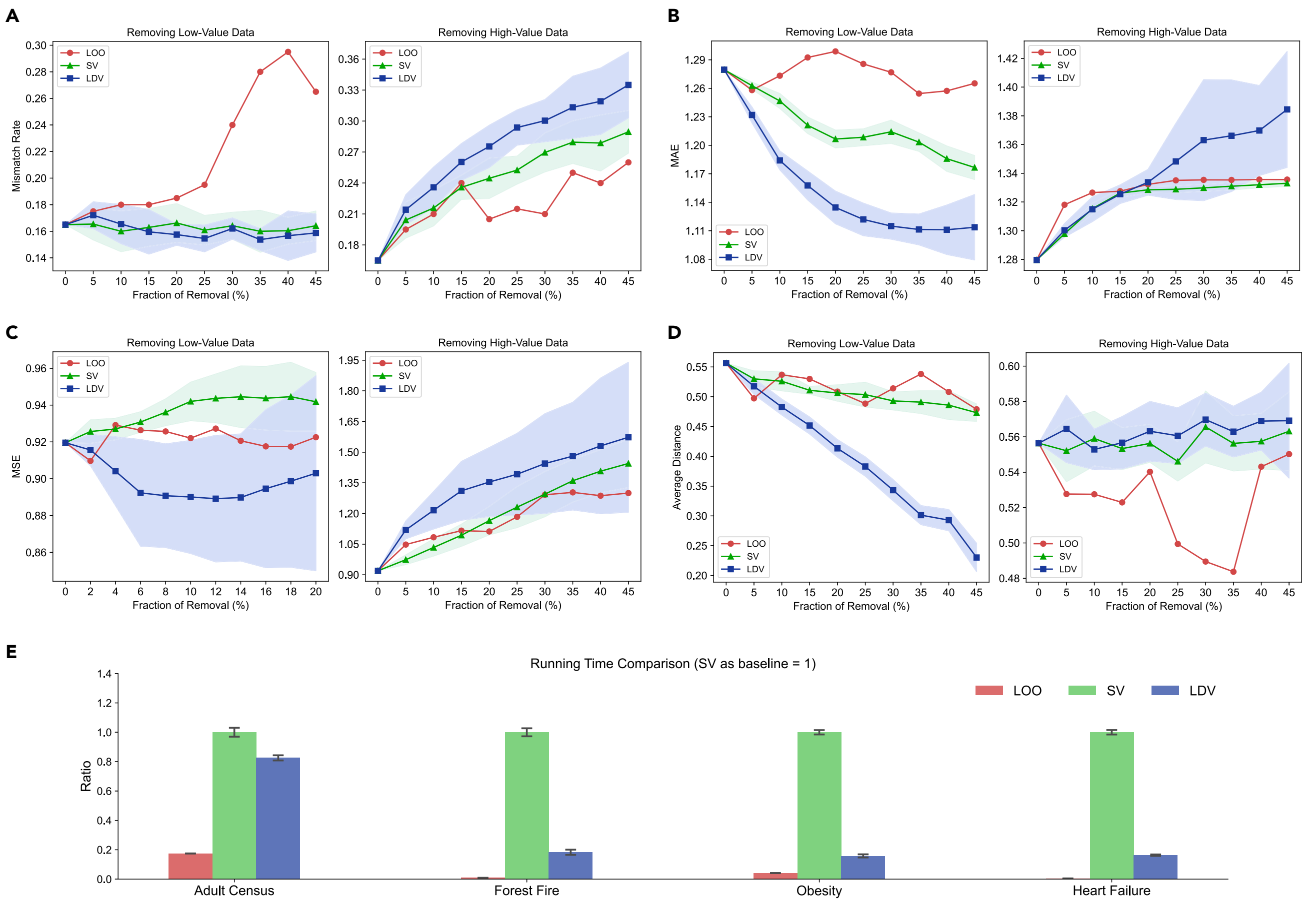

Figure 2. Comparative trends in high/low data removal across different methods

(A–D) This figure encapsulates the results from experiments conducted when using the leave-one-out (LOO) method, the Shapley value (SV), and the learning-based data-valuation (LDV) approach across four distinct datasets: (A) adult census, (B) forest fire, (C) obesity, and (D) heart failure. Each subgraph includes two representations of high/low-value removal. The depicted curves represent the mean values; the shaded areas around these curves indicate confidence intervals.
(E) The runtime ratio of various methods with runtime for the SV as the baseline = 1.

Applicability in wind power prediction

Wind power, heralded as a key player in the global energy shift, is at the forefront due to its sustainability, affordability, and abundant nature.[38,39] Acknowledging the inherent volatility and unpredictability of wind-energy generation, we applied predictive modeling as the data utilization case, and we selected accuracy as a metric for data-value assessment. Accurate wind-power forecasting mitigates uncertainty and facilitates optimized energy dispatch; this approach can yield substantial savings, thus augmenting the efficiency of electricity markets.[40,41] Consequently, the application of data-valuation methodologies is imperative for enhancing forecast precision by leveraging reliable datasets with informative, high-quality data points.[42–44]

Our predictive model inputs include wind-power data from the preceding 48 h coupled with daily and hourly meteorological information to forecast subsequent 24-h power generation. We adopted the mean absolute percentage error (MAPE) as a reverse metric to gauge the utility of the prediction model trained on various data subsets. In this section, we focus on data-valuation experiments utilizing wind-power data from Sichuan (SC) Province, which is renowned for its significant power generation and demand.

Figure 3A exhibits a data-value pattern consistent with earlier experiments, with the removal of high-value data correlating to a pronounced increase in the wind-power prediction error. Eliminating 40% of the high-value data inflates the median MAPE by 25.8%. Conversely, discarding low-value data reduces the MAPE, peaking at a 4.5% decrease with 25% data removal, equivalent to a 10.5% relative change. This finding suggests that excising superfluous data may indeed bolster the accuracy of renewable-energy forecasts. However, the subsequent rebound in the MAPE above the 25% removal threshold underscores the risk of depleting the training sample to a detrimental level; this depletion can undermine model training efficacy.

Figure 3B illustrates the impact of excising 20% of the valued data on the prediction error across a 24-h forecast period. The baseline error trajectory, derived from training with the complete dataset, pinpoints the primary forecasting challenges as occurring beyond the 10-h mark, where the temporal inertia of the power-generation time series starts to wane. The removal of 20% of the low-value data leads to a tangible reduction in prediction

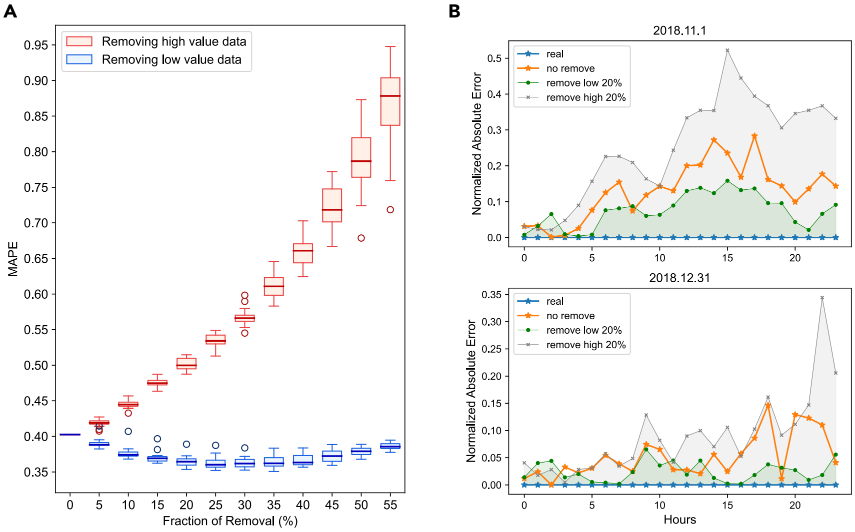

Figure 3. Value characteristics of wind-power data in SC

(A) The red and blue boxplots show the MAPE variation from removing the highest- and lowest-value data, respectively, with the median marked by the bold line.

(B) The error trends on the first and last days of the last 2 months from the test set (11.1 and 12.31) illustrate that the absolute prediction error changes when 20% of the highest-value (gray) and lowest-value (green) data are removed compared to that of the full dataset (orange).

error after 10 h, whereas the excision of an equal proportion of high-value data markedly exacerbates the error. Thus, the strategic value of wind-power data is demonstrated by their ability to confine prediction errors, particularly beyond the critical 10-h threshold. By eliminating low-value data, we reduce noise and enhance the training process of the model. In contrast, removing high-value data impairs the predictive performance of the model by eliminating essential training data. This delineation of results validates the practical application and robustness of our data-valuation methodology in wind-power forecasting, especially from an hourly perspective.

Intrinsicality and transferability in subset valuation

In this study, we investigated value patterns within wind-energy data subsets and their learning performance in terms of value. We split the initial 42.8% of the training set and the remaining portion into two datasets, named alpha and beta. Then, we trained two different learners: (1) by using both alpha and beta, and (2) by using only alpha. These learners were then tasked with evaluating both datasets. The learner trained on both datasets serves as the optimal benchmark and possesses comprehensive training information. We compared its data-value output with that of the learner trained solely on alpha, assessing consistency across both the trained (alpha) and untrained (beta) datasets.

Figure 4 reveals a significant correlation in the data-value outputs between the learners. An R^2 value of 0.81 for alpha reflects the stability and intrinsic robustness of our learning module within wind-energy scenarios. These facts indicate reliable value generalization for the same dataset under varying training conditions. Moreover, an R^2 value of 0.76 for beta highlights the learner's strong generalizability with evidence for the transferability of known value patterns to similar but previously unseen datasets.

The implications of such robustness and generalizability are 2-fold. First, consistent value patterns observed across different training regimens underscore a stable, intrinsic link between wind-power data and their value. This consistency reinforces the importance and necessity of our research into data valuation in the wind-energy domain. Second, the ability to accurately predict the value of new, untrained data demonstrates the transferability of value patterns. Given the extensive computational resources typically required for processing large volumes of wind-power data, the ability to generalize and to transfer learning efficiently is pivotal; these abilities potentially reduce the computational burden associated with extensive training.

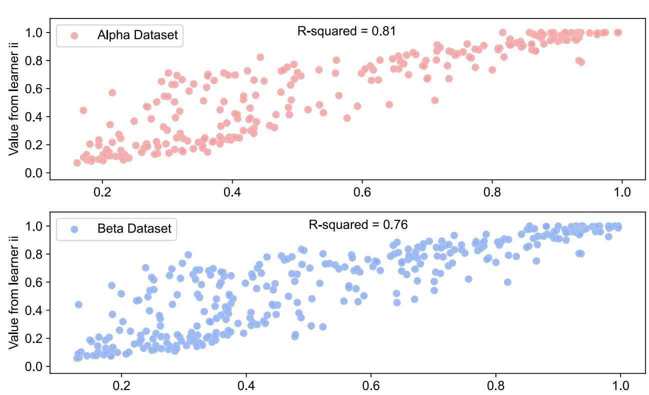

Figure 4. Data-value patterns between subsets

The graph illustrates the correlation in data values as assessed by learner i (trained on the alpha dataset and beta dataset) and learner ii (trained only on the alpha dataset), where red represents data points in (A) alpha and blue represents (B) beta. R-squared values are presented to quantify the linear correlation after fitting a linear regression model with ordinary least squares.

Geographic value patterns in China's wind power prediction

In our study, we extend the data-valuation methodology nationally to discern patterns across 25 Chinese provinces. Employing consistent data-value calculations, we integrate the effects of high/low-value data removal into Figure 5, utilizing the province abbreviations listed in Table S1.

The data-value behavior across most Chinese provinces exhibits a consistent pattern. In general, the MAPE decreases initially then increases with the sequential removal of low-value data, while the removal of high-value data leads to a steady increase in the MAPE. This trend is exemplified by Jilin (JL) and Hebei (HE), which exhibit curve patterns similar to those of SC. In addition, the varied geographic and climatic conditions across

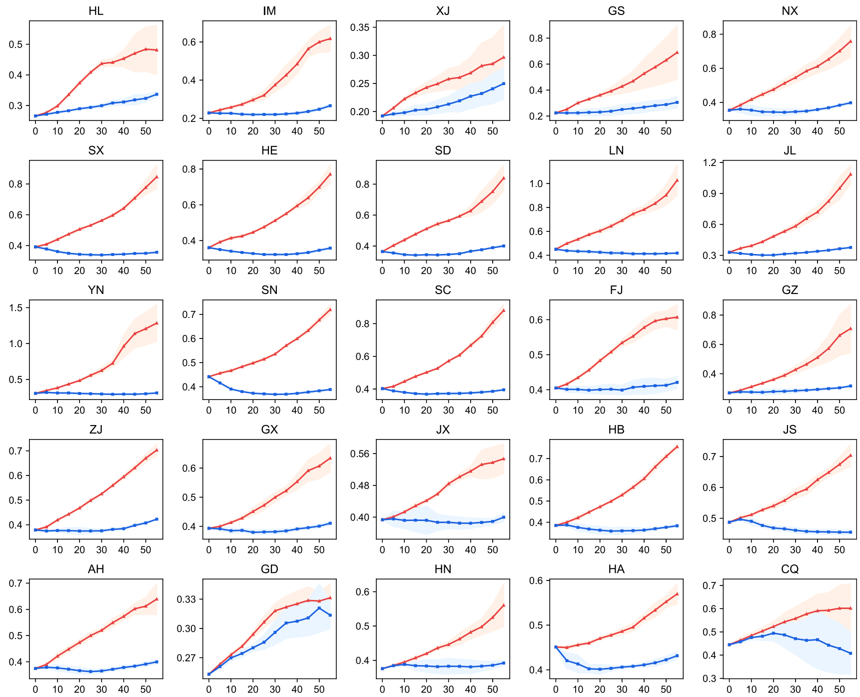

Figure 5. Data-value effects on wind-power prediction across 25 provinces in China
Prediction performance across 25 provinces relative to the removal of high-value (red) and low-value (blue) data. The vertical axis shows the MAPE; the horizontal axis represents the percentage of data removed.

China indicate that a low data value does not necessarily equate to model interference, as observed in provinces such as Heilongjiang (HL) and Guangdong (GD), where removing low-value data does not substantially decrease the MAPE. However, the universal increase in the MAPE with the removal of high-value data across all 25 provinces emphasizes the critical nature of certain feature data. The marked difference in the MAPE following the removal of equivalent amounts of high- and low-value data accentuates the capacity of the algorithm to discriminate data values. This consistency, despite China's geographic and climatic variability, attests to the adaptability and efficacy of the algorithm on distributed datasets; these benefits underscore the robustness and versatility of our approach.

The results reveal significant regional disparities in the sensitivity of wind-power prediction to data valuation across Chinese provinces, as illustrated by the extent of divergence in the curves for low- and high-value data removal. This sensitivity highlights the varying impacts on prediction accuracy resulting from the elimination of equal quantities of data with differing values. Provinces such as HE and Liaoning (LN) demonstrate pronounced increases in the MAPE difference following data removal, in contrast to the more subdued patterns from Xinjiang (XJ) and Chongqing (CQ). These observed variations in data-value sensitivity across provinces prompted further investigation into the underlying patterns and reasons for these geographic patterns.

In subsequent analysis, we focus on the distributed sensitivity of the data across provinces, with a cap of 40% on the removed data volume. Specific removal points at 10% intervals up to 40% provide insight into the predicted effects created by valued wind-power data. In Figures 6A, 6B, and 6C, color-coded deviations indicate the impact of data valuation on the prediction performance (MAPE) difference, with darker hues denoting greater sensitivity. An observable pattern emerges, with darker shades from southwest to northeast China occurring at 10% removal, indicating a sensitive region. This sensitivity becomes more defined within a narrower belt at 20% removal, shifting from northeast to southwest with increased removal (40%).

In Figure 6D, we calculate a generalized sensitivity index for each province by averaging the effect ranging from 10% removal to 40% removal; we identify a continuous geographic belt of heightened sensitivity, labeled S2 within the red boundary. This belt segregates the provinces into three distinct and continuous regions: S2 within the boundary, S1 to the northwest, and S3 to the southeast. This delineation aids in quantifying and in understanding the geographical distribution of data-valuation sensitivity in national wind-power forecasting.

In our investigation of geographic patterns of data sensitivity, we merge data valuation with energy dynamics data to elucidate regional differences across Chinese provinces. Figure 7A displays a scatterplot distinguishing provinces by sensitivity along two variance dimensions. A clear division among provinces in three regions (S1, S2, S3) is visible, with the high data-value sensitivity of S2 distinctively in the upper right. Notably, within S2, Gansu (GS) Province connects the climatic zones of XJ and Inner Mongolia (IM), which exhibit characteristics similar to those of the S1 region.

Horizontally, a broader data-value variance in S2 suggests an unstable system; moreover, the removal of extreme values significantly affects the distribution bounds, altering the central tendency and data spread. This phenomenon underscores how greater value variance enhances sensitivity to data removal, directly influencing model accuracy in predictive tasks, particularly in diverse datasets such as wind-power forecasting. Vertically, there is a notable difference in the wind-power variance between the S1 and S3 regions. The power variance of S1 is considerably greater than that of S3, with S2 falling in between.

For a detailed interpretation, Figure 7B presents the normalized power output distributions for these regions. The peak power metric indicates the most prevalent normalized output level within each region, embodying the percentile relative to the maximum observed output.

Influenced by the subtropical monsoon climate of the southeast plains of China, S3 has the lowest peak power at 15.43%, which aligns with the limited wind resources and consistent low-level outputs with low variance. Consequently, data valuation has a limited role in distinguishing performance due to the predictability of low, stable outputs. Conversely, the peak power of 34.76% in S1 signals substantial wind-energy production

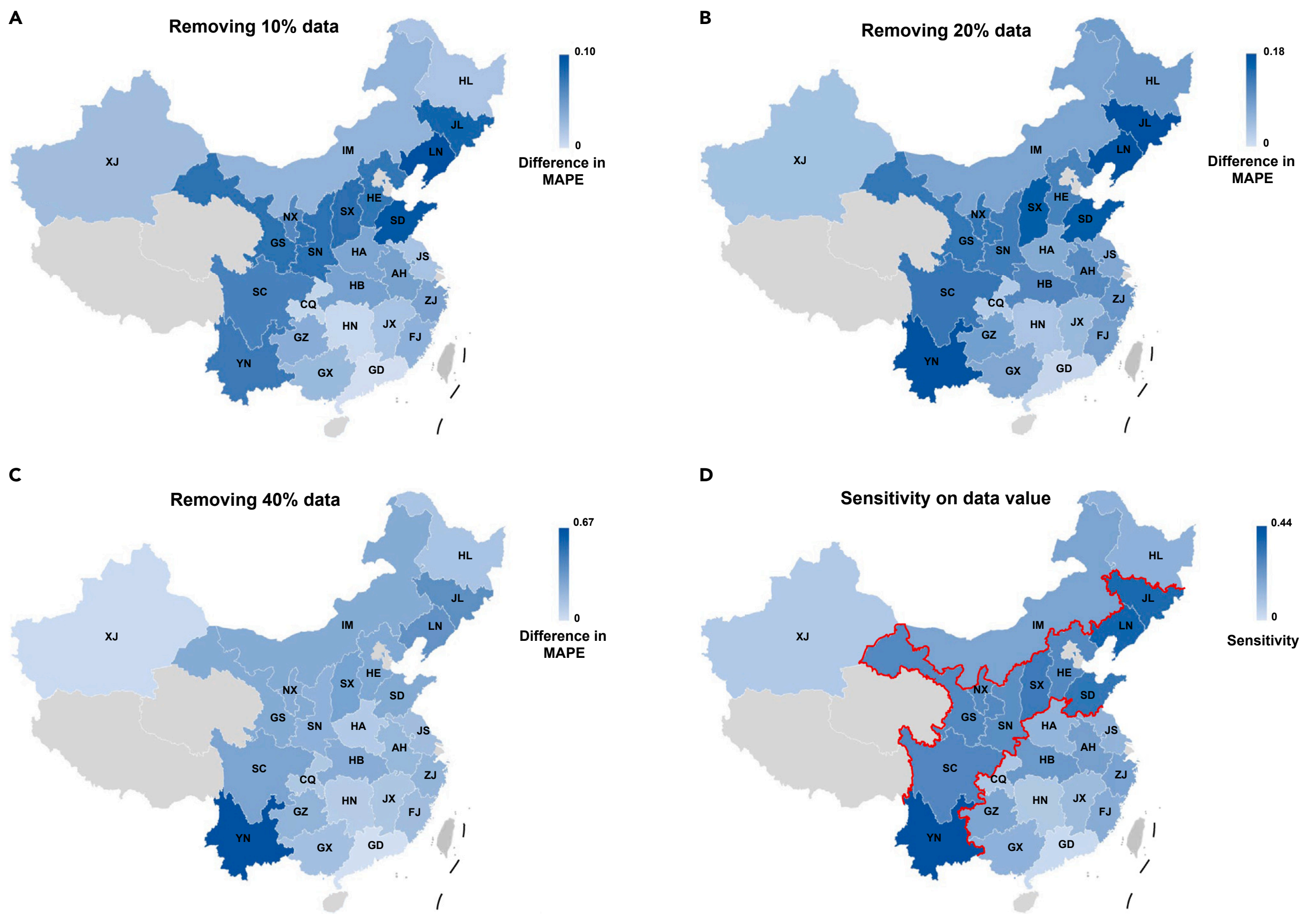

Figure 6. Sensitivity analysis of nationwide data values

(A–C) The sensitivity of the data valuation across provinces when (A) 10%, (B) 20%, and (C) 40% of the highest and lowest value data are removed, respectively, with color intensity indicating the extent of MAPE variation.

(D) The overall sensitivity of each province to the data value, with the color intensity reflecting the average impact of removing 10% of the data on model performance. The red line marks the region with the highest sensitivity (S2), dividing China into three zones: S1 (northwest), S2 (central belt), and S3 (southeast).

with notable variability, influenced by its temperate continental climate. Despite these fluctuations, the cyclical nature of the wind patterns in S1 simplifies predictions, leading to a homogenized data value and a reduced impact of data variations on model performance.

The "data-favored" profile of S2 emerges with its 26.65% peak power, situated between the stability of S3 and the variability of S1. Its location, encompassing both temperate and subtropical climates, results in a complex, less-predictable wind pattern. This variability necessitates precise data valuation for capturing the intricacies of wind generation, for underscoring the critical role of high-quality data, and for removing redundant data for accurate forecasting. This is particularly true for the three most sensitive provinces in S2—Yunnan (YN), LN, and JL—for which refined data valuation is essential for enhancing model efficiency and predictive reliability.

DISCUSSION

In this study, we present a learning-based paradigm for data valuation, and we demonstrate significant improvements in model utility and efficiency across various data-driven contexts. By employing deep reinforcement learning in our approach, we select the data subsets with the highest utility, optimizing performance, as shown by comparative analyses across diverse scientific datasets. Specifically, in wind-power prediction, our paradigm proficiently identifies valuable data, influencing model accuracy through the removal of data points and demonstrating consistent data-value patterns across regions.

Our findings reveal intrinsic and transferable data-value patterns, with learners effectively capturing value characteristics despite limited training data. This indicates the underlying nature of the data value and the potential for computational efficiency. Nationwide application of our algorithm has verified the presence of uniform data-value patterns across 25 Chinese provinces. Nonetheless, the degree to which the predictive accuracies of different regions react to data removal varies illuminates diverse spatial sensitivities to data valuation. This analysis leads to the identification of a geographically sensitive belt, providing a deeper understanding of the regional nuances in the value patterns of wind-power data.

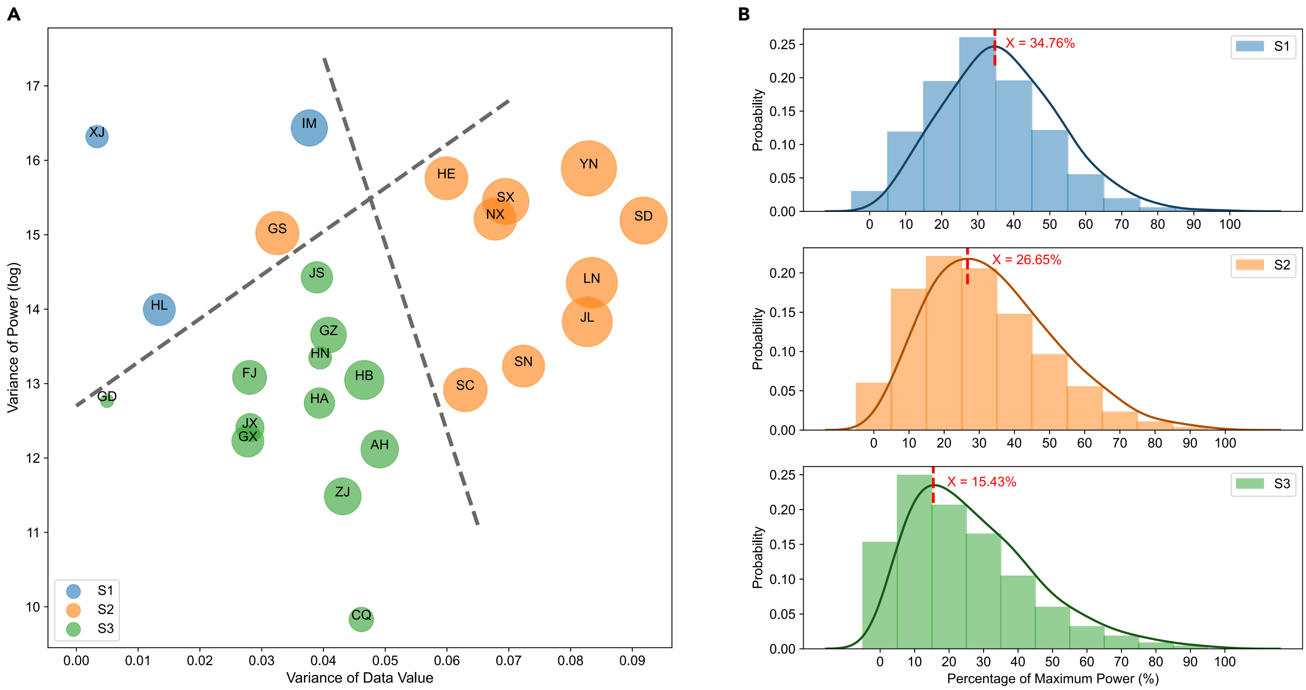

Figure 7. Explanation of the data-value-sensitive geographic belt formation

(A) Scatterplot for 25 Chinese provinces displaying data-value variance on the horizontal axis and wind-power-generation variance (logarithmic) on the vertical axis, with bubble size indicating sensitivity levels. Provinces are color coded: S1 (blue), S2 (orange), and S3 (green), demarcated by two dotted lines.

(B) Annual wind-power-generation distribution across the three regions, with power-generation percentiles on the horizontal axis and probability on the vertical axis. Decile histograms and kernel density estimation (KDE) curves illustrate distribution disparities.

Our data-valuation framework offers substantial contributions to renewable-energy policy by enhancing wind-power forecasting and by mitigating uncertainty in energy systems. We introduce an efficient approach for data-level governance, optimizing computational resources by effectively evaluating wind-power data quality and implementing strategic data filtering. This approach enhances forecast accuracy, stabilizes the energy grid, and supports potential data transactions.

The framework underscores the importance of targeted data collection and analysis, particularly in data-sensitive regions such as the S2 area in China. Here, the removal of specific valuable data markedly influences the predictive performance. Provinces within this belt are advised to advance their data collection and analytical capabilities, potentially through expanded sensor deployment and advanced processing techniques. Such measures are vital for effectively managing wind-power fluctuations and ensuring robust, field-applicable datasets.

By utilizing data-driven insights, provinces can refine their energy management strategies. The flexibility and comprehensive assessment of our framework can equip policy makers with crucial tools for crafting bespoke renewable-energy policies, thereby promoting a sustainable and economically beneficial energy ecosystem.

Future research directions can broaden the application of this paradigm to address a range of practical decision-making challenges. This includes optimizing power dispatch in the face of renewable-energy variability and leveraging enhanced data analysis to improve public health outcomes. Additionally, customizing data governance frameworks to align with specific regional energy policies may increase the effectiveness of data-driven decision making. Such an approach not only translates data value into tangible economic benefits but also bolsters efforts toward sustainable development goal.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

For further details, questions, or requests, direct correspondence to Jianxiao Wang at wang-jx@pku.edu.cn.

Materials availability

No physical materials were utilized in the conduct of this research.

Data and code availability

- The source codes of algorithms and datasets used in comparative experiments are available at Figshare: https://doi.org/10.6084/m9.figshare.24315535.[45]
- The wind-power and meteorological datasets are available at Figshare: https://doi.org/10.6084/m9.figshare.24316738.[46]
- The source codes of valuation in renewable uncertainty are available at Figshare: https://doi.org/10.6084/m9.figshare.24317092.[47]
- All experimental results presented in this article are openly accessible as data and are available at Figshare: https://doi.org/10.6084/m9.figshare.24320974.[48]

Methods

LOO and SV

The LOO method is a foundational technique in data valuation. It involves assessing the marginal value of an individual data point by comparing it to its complement, as follows:

$$L(i) = V(D) - V(D \backslash \{i\})$$

where $L(i)$ is the LOO value of an individual data point i, determined by quantifying the discrepancies in utility between the entire dataset $V(D)$ and its complement $V(D \backslash \{i\})$.

Stemming from efficient allocation of utility in game theory and a more comprehensive consideration of data contributions, the data SV is defined as follows:

$$Q(i) = \sum_{S \subseteq D \backslash \{i\}} \frac{|S|!(N - |S| - 1)!}{N!} (V(S \cup \{i\}) - V(S))$$

where $Q(i)$ denotes the SV of an individual data point i, treated as a participant in the game, is computed by taking the average of the marginal contributions from all the other data subsets S in the entire dataset D (with N data points) while considering the metric that defines the value based on the dataset. Considering its exponential computational complexity, the SV calculation in this study is approximated by using a Monte Carlo simulation with randomized sorting.[22] This method provides an efficient way to conduct an abbreviated computation, maintaining the integrity of the conceptual framework of the SV while addressing the computational demands.

LDV design

In our research, we introduce a data-valuation paradigm that leverages deep reinforcement learning to optimize the utility of the model. We denote the entire dataset as $\mathcal{D} = \{\mathbf{d}_i\}_{i=1}^{N} \sim \mathcal{P}$, where $\mathbf{d}_i \in \mathcal{R}^m$ is an individual data point characterized by an m-featured vector. The utility derived from a subset S is defined as $V(S)$. Subset S corresponds to a N-dimensional filter vector $\mathbf{s} \in \{0,1\}^N$, where $S = \{\mathbf{d}_i | s_i = 1, i = 1, 2 \ldots N\}$. For clarity, we express this as $V(\mathcal{D}, \mathbf{s})$ instead of $V(S)$.

The data-value learner, represented as $l_\varphi : \mathcal{D} \rightarrow [0,1]$, is optimized to predict the probability of the inclusion of each data element in subset S. This probability is interpreted as the value of the data point. A higher $l_\varphi(\mathbf{d_i})$, approaching 1, suggests a greater likelihood of the data point $\mathbf{d_i}$ being selected for model training, indicating its significant contribution and higher value in practical scenarios. Conversely, a value closer to 0 implies a higher probability of the data point being excluded from the final task or training. We define the filter as a binomial distribution filter $F_B : [0,1] \rightarrow \{0,1\}$, indicating whether $\mathbf{d_i}$ is selected ($F_B = 1$) or not ($F_B = 0$). Hence, $s_i = F_B(l_\varphi(\mathbf{d_i}))$.

Given the nondifferentiable nature of this filtering process, we transform the utility of the data subset into a "sampling signal." This necessitates adopting a reinforcement-learning framework for resolution. We model this as a decision process with a single step, where the entire dataset $\mathcal{D}$ represents the "state," the filter vector $\mathbf{s}$ denotes the "action," the utility V is the "reward" from a data-driven environment, and the learner's parameter φ is the strategy to be optimized. Then, the probability of the filter vector $\boldsymbol{s}$ being actioned based on $l_\varphi(\mathcal{D})$ is $\pi_\varphi(\mathcal{D}, \mathbf{s}) = \prod_{i=1}^{N} [l_\varphi(\mathbf{d}_i)^{s_i} \cdot (1 - l_\varphi(\mathbf{d}_i))^{1-s_i}]$, the result of the binomial distribution filter F_B. To maximize the utility of the selected subset, we formulate the optimization problem for training the learner as follows:

$$\max_{\varphi} \mathop{E}_{\mathbf{s} \sim \pi_\varphi(\mathcal{D}, \cdot)} [V(\mathcal{D}, \mathbf{s})]$$

Policy gradient and improvements

Our algorithm undergoes theoretical refinement to enhance its efficiency and robustness. We address the issue of the non-differentiability of $V(\mathcal{D}, \mathbf{s})$ with respect to φ by using the policy gradient method. This technique allows for the effective training of the learner under the reinforcement-learning framework. If $\mathcal{J}(\varphi)$ is denoted as the objective function, the gradient of strategy φ is expressed as follows:

$$\nabla_\varphi \mathcal{J}(\varphi) = \mathop{E}_{\mathbf{s} \sim \pi_\varphi(\mathcal{D}, \cdot)} [V(\mathcal{D}, \mathbf{s}) \cdot \nabla_\varphi \log (\pi_\varphi(\mathcal{D}, \mathbf{s}))]$$

To enhance the stability of the learning process based on policy gradients, we employ a moving average of the previous rewards, denoted as δ, with a window size T (number of iterations) as the baseline.

Our analysis reveals that, without specific intervention, the values of all the data points ($\forall\ \mathbf{d_i} \in \mathcal{D}$) are highly likely to skew toward either all 0 s or all 1 s. Such extrema in the value distribution leads to suboptimal results. This local convergence is due to inadequate exploration of alternative optimum strategies. To counteract this, we introduce a penalty function $p(\varphi; \sigma, \boldsymbol{t})$ to promote exploration:

$$p(\varphi; \sigma, \mathbf{t}) = \sigma \cdot \left[max\left(\sum_{i=1}^{N} l_\varphi(\mathbf{d}_i) - N \cdot \mathbf{t}, 0 \right) + max\left(1 - \sum_{i=1}^{N} l_\varphi(\mathbf{d}_i) - N \cdot \mathbf{t}, 0 \right) \right]$$

where σ is the penalty factor and $\boldsymbol{t}$ is the threshold near 1.

To further enhance the stability, we propose using repeated sampling to maximize the sample utility. We incorporate importance sampling and a target network within the reinforcement-learning framework, facilitating a transition from on-policy to off-policy methods. The modified objective function is

$$\mathop{E}_{\mathbf{s} \sim \pi_{\hat{\varphi}}(\mathcal{D}, \cdot)} \left[\frac{\pi_\varphi(\mathcal{D}, \mathbf{s})}{\pi_{\hat{\varphi}}(\mathcal{D}, \mathbf{s})} V(\mathcal{D}, \mathbf{s}) \right]$$

Given the increased variance from off-policy sampling and the potential limitation of updates to new policies, we employ a clipping function to moderate the disparity between old and new policies:

$$clip\left(\frac{\pi_\varphi(\mathcal{D}, \mathbf{s})}{\pi_{\hat{\varphi}}(\mathcal{D}, \mathbf{s})}, 1 - \eta, 1 + \eta \right)$$

Here, the clip function $clip(x, a, b)$ truncates x the bounds a and b, maintaining the sensitivity of the policy gradient to larger step sizes within manageable limits η. Ultimately, we integrate the clipped objective with the penalty term to form the final objective function, guiding the training of the data learner's strategy via the policy gradient method.

Comparison experiments

The datasets utilized for the comparison experiments in our study were sourced from publicly available repositories comprising adult,[34] forest-fire,[35] obesity,[36] and heart-failure[37] data. Notably, in the heart-failure clustering analysis, we adopted an unsupervised learning approach, and we defined value by calculating the distance from each point to its nearest cluster center; this approach demonstrates the versatility of our framework in handling nonpredictive tasks.

To construct classification, regression, and clustering models, we used appropriate machine-learning algorithms tailored to each problem type. For classification tasks, we employed eXtreme Gradient Boosting (XGB), a sophisticated ensemble boosting model renowned for its effectiveness in featured data management, numerical pattern extraction, and predictive analysis capabilities.[49] Support vector regression (SVR) was chosen for the regression tasks, and another regression task utilized the light gradient-boosting machine (LGBM), a decision-tree-based gradient-boosting framework noted for its outstanding performance and scalability.[50] K-means clustering was applied for clustering analyses. In each iteration, the models were trained on a selected subset of the data based on the calculated value metric.

The inverse metric employed in classification is the "mismatch rate," defined as the complement of accuracy ($1 - accuracy$). For clustering, we computed the metric by aggregating the distances between each data point and its respective cluster center. As the volume of data is reduced, the cumulative distance decreases proportionally. To highlight the unique approach of our algorithm in data valuation, the vertical axis in Figure 2D represents the averaged distance, calculated as the sum of distances divided by the data volume. In our analysis, larger metric values indicate greater overall utility in the dataset. For each dataset, 12 training repetitions were performed for each method, and the results were averaged and are presented in Table 1.

Wind power prediction

We analyzed 17,520 h of wind data from 2017 to 2018, focusing on 25 Chinese provinces and regions and omitting those with insufficient or corrupt data.[51] This subset still covers a significant portion of China, excluding Hong Kong, Macau, Tianjin, Taiwan, Tibet, Qinghai, Beijing, Shanghai, and Hainan. To aid in wind-power forecasting, we integrated multidimensional numerical weather prediction (NWP) data from The National Aeronautics and Space Administration (NASA), encompassing more than 3,300 counties nationwide. Employing cluster averaging, we extracted characteristic NWP data for each

CellPress
OPEN ACCESS

province and compiled hourly and daily data to support wind-power predictions for corresponding hours. The predictive model inputs merge these NWP data with historical power-generation data spanning the previous 48 h, setting the stage to forecast hourly power generation for the subsequent 24 h. Leveraging the capabilities of the LGBM[50] for its proven effectiveness in time-series predictions, we ensured a deterministic and stable utility for power forecasting.[52]

The MAPE for the prediction model is calculated as follows:

$$\mathrm{MAPE} = \frac{100\%}{n * 24} \sum_{i=1}^{n} \sum_{t=1}^{24} \left| \frac{P_{i,t} - A_{i,t}}{max(A_{i,t}, \epsilon)} \right|$$

where n is the test set size. $P_{i,t}$ and $A_{i,t}$ denote the predicted and actual wind-power values, respectively, at hour t during the day. We introduced $\epsilon = 0.12$ in the MAPE calculation to prevent small fluctuations in low actual power outputs from unduly magnifying the error metric, thus avoiding an overemphasis on prediction accuracy during periods of minimal generation.

In the data-value learner model, data points are represented by input vectors x consisting of hourly/daily NWP and the preceding 48-h power data, with y representing the subsequent 24-h power data. These inputs are processed through a fully connected neural layer, and a sigmoid activation function computes the data value between 0 and 1:

$$Sigmoid(x) = \frac{1}{1+e^{-x}}$$

Sensitivity calculation for each province

To assess the sensitivity of wind-power forecasting to data removal across various provinces, we define Δ as a metric quantifying the average change in the MAPE per 10% increase in data removal, up to a limit of 40%. The formula for Δ is given by

$$\Delta = \delta_{10\%} + \frac{\delta_{20\%}}{2} + \frac{\delta_{30\%}}{3} + \frac{\delta_{40\%}}{4}$$

Here, $\delta_{m\%}$ denotes the change in MAPE resulting from the removal of m% of the highest or lowest value of the data. This methodology allows for a comprehensive assessment of the impact of data valuation in removal experiments, and this approach allows us to effectively quantify how data value influences wind-power prediction. Additionally, it serves as an aggregate indicator to evaluate the sensitivity of each province to the data value in relation to the prediction accuracy.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.patter.2024.100965.

ACKNOWLEDGMENTS

This work was supported by the National Key Research and Development Program of China (2022YFB2405600) and the National Natural Science Foundation of China (72241420, 52277092).

AUTHOR CONTRIBUTIONS

Y.W., J.W., and J.S. conceived and designed the research. Y.W. and J.W. developed the framework and formulated the theoretical model. Y.W., J.W., and F.G. carried out the data search. Y.W., J.W., and F.G. carried out the simulations. Y.W. and J.W. conducted the prediction analysis. All authors contributed to the discussions on the method and the writing of this article.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 3, 2023
Revised: February 3, 2024
Accepted: March 6, 2024
Published: March 27, 2024

REFERENCES

1. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., and Prabhat, f. (2019). Deep learning and process understanding for data-driven earth system science. Nature *566*, 195–204. https://doi.org/10.1038/s41586-019-0912-1.
2. Carleo, G., Cirac, I., Cranmer, K., Daudet, L., Schuld, M., Tishby, N., Vogt-Maranto, L., and Zdeborová, L. (2019). Machine learning and the physical sciences. Rev. Mod. Phys. *91*, 045002. https://doi.org/10.1103/RevModPhys.91.045002.
3. Wu, F., Lu, C., Zhu, M., Chen, H., Zhu, J., Yu, K., Li, L., Li, M., Chen, Q., Li, X., et al. (2020). Towards a new generation of artificial intelligence in China. Nat. Mach. Intell. *2*, 312–316. https://doi.org/10.1038/s42256-020-0183-4.
4. Oztemel, E., and Gursev, S. (2020). Literature review of industry 4.0 and related technologies. J. Intell. Manuf. *31*, 127–182. https://doi.org/10.1007/s10845-018-1433-8.
5. Shahbaz, M., Wang, J., Dong, K., and Zhao, J. (2022). The impact of digital economy on energy transition across the globe: The mediating role of government governance. Renew. Sustain. Energy Rev. *166*, 112620. https://doi.org/10.1016/j.rser.2022.112620.
6. Duch-Brown, N., and Rossetti, F. (2020). Digital platforms across the European regional energy markets. Energy Pol. *144*, 111612. https://doi.org/10.1016/j.enpol.2020.111612.
7. Adefarati, T., and Bansal, R.C. (2019). Reliability, economic and environmental analysis of a microgrid system in the presence of renewable energy resources. Appl. Energy *236*, 1089–1114. https://doi.org/10.1016/j.apenergy.2018.12.050.
8. Taleb, I., Serhani, M.A., Bouhaddioui, C., and Dssouli, R. (2021). Big data quality framework: a holistic approach to continuous quality management. J. Big Data *8*. 76-41. https://doi.org/10.1186/s40537-021-00468-0.
9. Bernhardt, M., Castro, D.C., Tanno, R., Schwaighofer, A., Tezcan, K.C., Monteiro, M., Bannur, S., Lungren, M.P., Nori, A., Glocker, B., et al. (2022). Active label cleaning for improved dataset quality under resource constraints. Nat. Commun. *13*, 1161. https://doi.org/10.1038/s41467-022-28818-3.
10. Jain, A., Patel, H., Nagalapatti, L., Gupta, N., Mehta, S., Guttula, S., Mujumdar, S., Afzal, S., Sharma Mittal, R., and Munigala, V. (2020). Overview and importance of data quality for machine learning tasks. In Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining, pp. 3561–3562. https://doi.org/10.1145/3394486.3406477.
11. Ghasemaghaei, M., and Calic, G. (2019). Can big data improve firm decision quality? The role of data quality and data diagnosticity. Decis. Support Syst. *120*, 38–49. https://doi.org/10.1016/j.dss.2019.03.008.
12. Emeis, S. (2018). Wind Energy Meteorology: Atmospheric Physics for Wind Power Generation (Springer).
13. Zhang, Y., Li, Y., and Zhang, G. (2020). Short-term wind power forecasting approach based on seq2seq model using nwp data. Energy *213*, 118371. https://doi.org/10.1016/j.energy.2020.118371.
14. Orlov, A., Sillmann, J., and Vigo, I. (2020). Better seasonal forecasts for the renewable energy industry. Nat. Energy *5*, 108–110. https://doi.org/10.1038/s41560-020-0561-5.
15. Liang, W., Tadesse, G.A., Ho, D., Fei-Fei, L., Zaharia, M., Zhang, C., and Zou, J. (2022). Advances, challenges and opportunities in creating data for trustworthy AI. Nat. Mach. Intell. *4*, 669–677. https://doi.org/10.1038/s42256-022-00516-1.
16. Triguero, I., García-Gil, D., Maillo, J., Luengo, J., García, S., and Herrera, F. (2019). Transforming big data into smart data: An insight on the use of the k-nearest neighbors algorithm to obtain quality data. Wiley Interdisciplinary Reviews: Data Min. Knowl. Discov. *9*, e1289. https://doi.org/10.1002/widm.1289.

17. Cichy, C., and Rass, S. (2019). An overview of data quality frameworks. IEEE Access *7*, 24634–24648. https://doi.org/10.1109/ACCESS.2019.2899751.
18. Paullada, A., Raji, I.D., Bender, E.M., Denton, E., and Hanna, A. (2021). Data and its (dis) contents: A survey of dataset development and use in machine learning research. Patterns *2*. https://doi.org/10.1016/j.patter.2021.100336.
19. Blázquez-García, A., Conde, A., Mori, U., and Lozano, J.A. (2021). A review on outlier/anomaly detection in time series data. ACM Comput. Surv. *54*, 1–33. https://doi.org/10.1145/3444690.
20. Yu, M., Wang, J., Yan, J., Chen, L., Yu, Y., Li, G., and Zhou, M. (2022). Pricing information in smart grids: A quality-based data valuation paradigm. IEEE Trans. Smart Grid *13*, 3735–3747. https://doi.org/10.1109/TSG.2022.3171665.
21. Wang, B., Guo, Q., Yang, T., Xu, L., and Sun, H. (2021). Data valuation for decision-making with uncertainty in energy transactions: A case of the two-settlement market system. Appl. Energy *288*, 116643. https://doi.org/10.1016/j.apenergy.2021.116643.
22. Ghorbani, A., and Zou, J. (2019). Data shapley: Equitable valuation of data for machine learning. In International conference on machine learning (PMLR), pp. 2242–2251. https://doi.org/10.48550/arXiv.1904.02868.
23. Jia, R., Dao, D., Wang, B., Hubis, F.A., Gurel, N.M., Li, B., Zhang, C., Spanos, C.J., and Song, D. (2019). Efficient task-specific data valuation for nearest neighbor algorithms. Preprint at arXiv. https://doi.org/10.48550/arXiv.1908.08619.
24. Yoon, J., Arik, S., and Pfister, T. (2020). Data valuation using reinforcement learning. In International Conference on Machine Learning (PMLR), pp. 10842–10851. https://doi.org/10.48550/arXiv.1909.11671.
25. Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M.C., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Noble, J., Barratt, D.C., and Hu, Y. (2022). Image quality assessment for machine learning tasks using meta-reinforcement learning. Med. Image Anal. *78*, 102427. https://doi.org/10.1016/j.media.2022.102427.
26. Wang, H., Lei, Z., Zhang, X., Zhou, B., and Peng, J. (2019). A review of deep learning for renewable energy forecasting. Energy Convers. Manag. *198*, 111799. https://doi.org/10.1016/j.enconman.2019.111799.
27. Elmachtoub, A., and Grigas, P. (2022). Smart "predict, then optimize". Manag. Sci. *68*, 9–26. https://doi.org/10.1287/mnsc.2020.3922.
28. Mišić, V.V., and Perakis, G. (2020). Data analytics in operations management: A review. Manuf. Serv. Oper. Manag. *22*, 158–169. https://doi.org/10.1287/msom.2019.0805.
29. François-Lavet, V., Henderson, P., Islam, R., Bellemare, M.G., and Pineau, J. (2018). An introduction to deep reinforcement learning. Foundations and Trends® in Machine Learning *11*, 219–354. https://doi.org/10.1561/2200000071.
30. Zhao, T., Hachiya, H., Niu, G., and Sugiyama, M. (2012). Analysis and improvement of policy gradient estimation. Neural Network. *26*, 118–129. https://doi.org/10.1016/j.neunet.2011.09.005.
31. Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimiza- tion. International conference on machine learning. PMLR, 1889–1897. http://dl.acm.org/doi/10.5555/3045118.3045319.
32. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. Preprint at arXiv. https://doi.org/10.48550/arXiv.1707.06347.
33. Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., and Madry, A. (2020). Implementation matters in deep policy gradients: A case study on ppo and trpo. Preprint at arXiv. https://doi.org/10.48550/arXiv.2005.12729.
34. Becker, B., and Kohavi, R. (1996). Adult (UCI Machine Learning Repository). https://doi.org/10.24432/C5XW20.
35. Cortez, P., and Morais, A.d.J.R. (2007). A data mining approach to predict forest fires using meteorological data. https:/api.semanticscholar.org/CorpusID:36868619.
36. Palechor, F.M., and Manotas, A.d.l.H. (2019). Dataset for estimation of obesity levels based on eating habits and physical condition in individuals from colombia, peru and mexico. Data Brief *25*, 104344. https://doi.org/10.1016/j.dib.2019.104344.
37. Chicco, D., and Jurman, G. (2020). Machine learning can predict survival of patients with heart failure from serum creatinine and ejection fraction alone. BMC Med. Inf. Decis. Making *20*, 16. https://doi.org/10.1186/s12911-020-1023-5.
38. Wang, Z., and Liu, W. (2021). Wind energy potential assessment based on wind speed, its direction and power data. Sci. Rep. *11*, 16879. https://doi.org/10.1038/s41598-021-96376-7.
39. Vargas, S.A., Esteves, G.R.T., Maçaira, P.M., Bastos, B.Q., Cyrino Oliveira, F.L., and Souza, R.C. (2019). Wind power generation: A review and a research agenda. J. Clean. Prod. *218*, 850–870. https://doi.org/10.1016/j.jclepro.2019.02.015.
40. Yan, J., Liu, Y., Han, S., Wang, Y., and Feng, S. (2015). Reviews on uncertainty analysis of wind power forecasting. Renew. Sustain. Energy Rev. *52*, 1322–1330. https://doi.org/10.1016/j.rser.2015.07.197.
41. Wang, Y., Zou, R., Liu, F., Zhang, L., and Liu, Q. (2021). A review of wind speed and wind power forecasting with deep neural networks. Appl. Energy *304*, 117766. https://doi.org/10.1016/j.apenergy.2021.117766.
42. Aryandoust, A., Patt, A., and Pfenninger, S. (2022). Enhanced spatio-temporal electric load forecasts using less data with active deep learning. Nat. Mach. Intell. *4*, 977–991. https://doi.org/10.1038/s42256-022-00552-x.
43. Wang, Y., Hu, Q., Li, L., Foley, A.M., and Srinivasan, D. (2019). Approaches to wind power curve modeling: A review and discussion. Renew. Sustain. Energy Rev. *116*, 109422. https://doi.org/10.1016/j.rser.2019.109422.
44. Zheng, L., Hu, W., and Min, Y. (2015). Raw wind data preprocessing: A data-mining approach. IEEE Trans. Sustain. Energy *6*, 11–19. https://doi.org/10.1109/TSTE.2014.2355837.
45. Wang, Y., Wang, J., Gao, F., and Song, J. (2023). A Learning-Based Data Valuation Method: Comparative Experiments (figshare). https://doi.org/10.6084/m9.figshare.24315535.
46. Wang, Y., Wang, J., Gao, F., and Song, J. (2023). Wind Power and Meteorological Data for 25 Provinces in china (2017-2018) (figshare). https://doi.org/10.6084/m9.figshare.24316738.
47. Wang, Y., Wang, J., Gao, F., and Song, J. (2023). Nationwide data valuation with renewable uncertainty analysis (figshare). https://doi.org/10.6084/m9.figshare.24317092.
48. Wang, Y., Wang, J., Gao, F., and Song, J. (2023). Data of figures in "unveiling value patterns via deep reinforcement learning in heterogeneous data analytics" (figshare). https://doi.org/10.6084/m9.figshare.24320974.
49. Chen, T., and Guestrin, C.X. (2016). A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pp. 785–794. https://doi.org/10.1145/2939672.2939785.
50. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. In Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 3149–3157. http://dl.acm.org/doi/abs/10.5555/3294996.3295074.
51. Lu, X., McElroy, M.B., Peng, W., Liu, S., Nielsen, C.P., and Wang, H. (2016). Challenges faced by china compared with the us in developing wind power. Nat. Energy *1*, 16061–16066. https://doi.org/10.1038/nenergy.2016.61.
52. Sun, X., Liu, M., and Sima, Z. (2020). A novel cryptocurrency price trend forecasting model based on lightgbm. Finance Res. Lett. *32*, 101084. https://doi.org/10.1016/j.frl.2018.12.032.