

Quality over Quantity: Boosting Data Efficiency Through Ensembled Multimodal Data Curation

Jinda Xu^{1*}, Yuhao Song^{2*}, Daming Wang², Weiwei Zhao¹, Minghua Chen²,
Kangliang Chen^{2†}, Qinya Li^{1‡}

¹Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University

²HAOMO.AI Technology Co., Ltd.

daidan2502@sjtu.edu.cn, {songyuhao, wangdaming}@haomo.ai, zhaoweimei@sjtu.edu.cn, {chenminghua, chenkanliang}@haomo.ai, qinyali@sjtu.edu.cn

Abstract

In an era overwhelmed by vast amounts of data, the effective curation of web-crawl datasets is essential for optimizing model performance. This paper tackles the challenges associated with the unstructured and heterogeneous nature of such datasets. Traditional heuristic curation methods often inadequately capture complex features, resulting in biases and the exclusion of relevant data. We introduce an advanced, learning-driven approach, **Ensemble Curation Of DATA ThroUgh Multimodal Operators (EcoDatum)**, incorporating a novel quality-guided deduplication method to ensure balanced feature distributions. EcoDatum strategically integrates various unimodal and multimodal data curation operators within a weak supervision ensemble framework, utilizing automated optimization to score each data point effectively. EcoDatum, which significantly improves the data curation quality and efficiency, outperforms existing state-of-the-art (SOTA) techniques, ranked 1st on the DataComp leaderboard, with an average performance score of 0.182 across 38 diverse evaluation datasets. This represents a 28% improvement over the DataComp baseline method, demonstrating its effectiveness in improving dataset curation and model training efficiency.

Introduction

The vast amount of data presents opportunities for training advanced deep learning models, but it also introduces significant noise and irrelevant information, which can hinder model effectiveness. In both academia and industry, the need for robust data curation techniques to extract meaningful signals from extensive digital information has become a pressing concern. In web-crawled datasets, data curation is a multi-faceted task involving various stages and methodologies. The core objective is to identify and retain high-quality samples while discarding noise or mitigating the impact of irrelevant data (Hoffmann et al. 2022). This process is crucial for optimizing model performance in the deep learning framework.

*These authors contributed equally.

†Project leader.

‡Corresponding author.

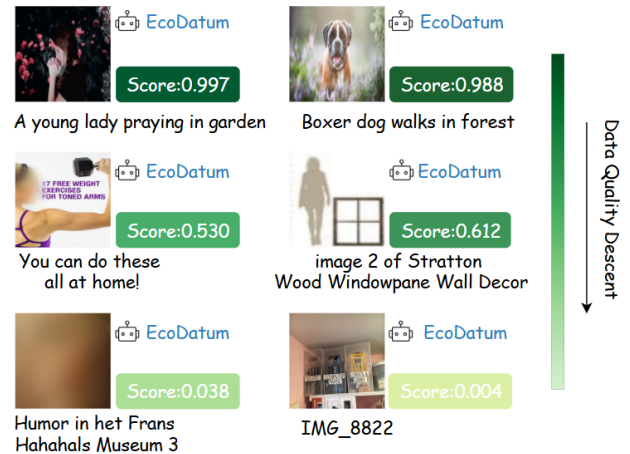


Figure 1: Web-crawled image-text datasets often vary in quality. EcoDatum addresses this by integrating unimodal and multimodal operators to discriminate data quality, guiding the selection process to enhance data curation efficiency.

Web-crawled data is inherently unstructured, diverse, and constantly evolving, making it essential to develop adaptive curation methods capable of handling such complexity (Hoffmann et al. 2022). Traditionally, data curation methods have relied heavily on heuristic filtering approaches based on manually defined content attributes, such as image resolution, graphic geometry, textual length, and linguistic complexity. While these heuristic methods provide a basic means of recognizing low-quality samples, they fail to adequately capture the subtle features of web-crawled data and may introduce biases or overlook relevant information (Mindermann et al. 2022; Maini, Yaghini, and Papernot 2021). To address these limitations, researchers are increasingly adopting automated curation methods that leverage deep learning techniques, including natural language processing, computer vision and cross-modal representation learning, to achieve a balance of quality and quantity (Schuhmann et al. 2021; Brown 2020; Torkashvand, Jameii, and Reza 2023).

This research proposes a data curation framework called EcoDatum to address the aforementioned issues. Specifi-

cally, we implement a range of efficient data curation strategies as operators, to enhance the data curation process and to achieve cross-modal data alignment at various levels of granularity, as detailed in Figure 1. However, a simple combination of these operators may introduce bias and lead to insufficient utilization of their individual strengths. To fully capture their synergies, we develop a weak supervision ensemble framework that integrates these operators, achieving a synergistic effect. Furthermore, to enhance the integration efficiency of unimodal and multimodal operators, EcoDatum introduces an automated optimization approach. This is achieved by tuning the weak supervision integration module using a composite metric and a tiny-labeled dataset.

As a novel weak supervision-based framework for multimodal data curation, EcoDatum achieves state-of-the-art performance on the DataComp data filtering track (Gadre et al. 2024). The visual language model trained on curated data demonstrates outstanding results over 38 downstream tasks, highlighting its strong generalizability. Extended experiments demonstrate the effectiveness of this research in understanding various operators in cross-modal data management, offering insights for future work.

Our main contributions are as follows:

1. We propose an auto-optimized ensemble framework, EcoDatum, which integrates techniques to enhance data quality and curate multimodal data, ensuring aligned, high-quality inputs for visual language pretraining.
2. We introduce a search-based optimization algorithm for weak supervision labeling function tuning that enhances the curation process and boosts the system’s robustness.
3. EcoDatum surpasses existing state-of-the-art techniques in the DataComp benchmark over 38 downstream tasks and ranks 1st on the leaderboard¹.

Related Work

Data Curation for Web-crawled Datasets

Recent research underscores the critical role of data curation in enhancing model performance with large-scale image-text datasets. Various studies focus on improving dataset quality through curation methods, such as enhancing the descriptiveness and cross-modal feature alignment of image-text pairs, and reducing redundancy (Radenovic et al. 2023; Nguyen et al. 2024; Abbas et al. 2023).

In a broader context, DataComp is a benchmark designed to evaluate the performance of multimodal models on large-scale, real-world datasets (Gadre et al. 2023). Recent advancements in the DataComp benchmark highlight notable progress in data curation techniques. Yokoo et al. (Yokoo et al. 2023) advanced data filtering using image-text similarity and caption modification, achieving notable progress in the Filtering and Bring Your Own Data (BYOD) Tracks. Yu et al. (Yu et al. 2023) evaluated data selection criteria’s impact on model performance, while Chen et al. (Chen et al. 2024) introduced DataJuicer for managing large-scale

datasets. Nguyen et al. (Nguyen et al. 2024) enhanced image captioning for better modality alignment, and Maini et al. (Maini et al. 2023) presented T-MARS for improved visual representation by bypassing text feature learning. Additionally, significant contributions have been made in the areas of synthetic data, contrastive learning, image-text alignment, and few-shot learning (Chen et al. 2020; Radford et al. 2021; Jia et al. 2021; Li et al. 2022; Alayrac et al. 2022).

Despite the significant advancements in data curation achieved by Radenovic et al. (Radenovic et al. 2023) and Nguyen et al. (Nguyen et al. 2024) through enhancing data relevance, existing automated filtering methods may still exclude valuable but less conventional data points or introduce biases by focusing too narrowly on specific aspects, potentially overlooking broader contextual information.

Ensemble Learning

Ensemble learning, which combines multiple models to improve performance and generalization, has long been a foundational approach in machine learning. Classic methods such as Bagging (Breiman 1996) and Boosting (Freund and Schapire 1997) initially demonstrated how model aggregation could reduce variance and improve accuracy.

Afterwards, ensemble learning has been increasingly applied to specialized tasks. Zimek, Schubert, & Kriegel (Zimek, Schubert, and Kriegel 2012) highlighted how ensemble methods enhance outlier detection by aggregating results from multiple models. Beluch et al. (Beluch et al. 2018) demonstrated that ensemble-based uncertainty sampling can significantly improve efficiency in active learning. Rasmus et al. (Rasmus et al. 2015) demonstrated that ensemble techniques enhance semi-supervised learning by effectively utilizing both labeled and unlabeled data. Song et al. (Song et al. 2022) used ensemble methods to improve data quality by detecting and filtering noisy or mislabeled data.

These advancements illustrate how ensemble techniques refine data preprocessing and improve model inputs.

Method

Overview

As illustrated in Figure 2, EcoDatum enhances the pretraining effectiveness of multimodal models like CLIP (Radford et al. 2021) by strategically selecting high-quality subset \hat{S} from the original dataset S . This targeted data curation improves the model’s zero-shot performance on diverse tasks. The framework utilizes an ensemble of specialized data operators for comprehensive quality assessment, which addresses various dimensions including image filtering, text analysis, and cross-modal alignment at multiple granular levels. Automated optimization enables the weak supervision system to generate quality scores for data samples, thus minimizing manual input and enhancing the precision of threshold settings. Consequently, EcoDatum streamlines the data curation process and significantly elevates the quality, ensuring the dataset meets the rigorous requirements for model training.

¹<https://www.datacomp.ai/dcclip/leaderboard.html>

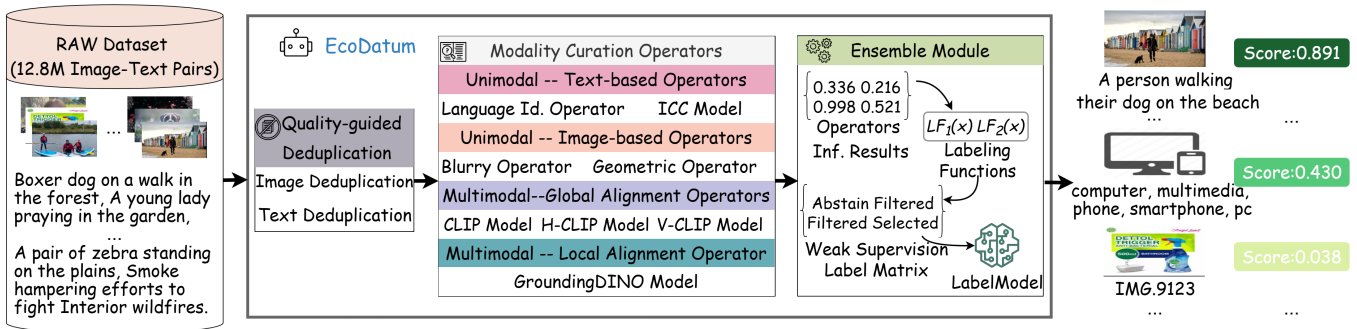


Figure 2: Overview of the EcoDatum Framework. EcoDatum utilizes quality-guided deduplication along with an ensemble of unimodal and multimodal data curation operators, that strategically curate multimodal datasets. This integrated approach systematically scores each data point, ensuring optimal quality and alignment for effective visual-language pretraining.

Quality Guided Deduplication

To improve our dataset’s diversity and distribution, we employ a quality-guided deduplication process that removes redundant text-image pairs. This approach uses perceptual hashing (Farid 2021) to generate hash codes, identifying duplicates based on visual and textual content. Subsequently, the CLIP model assesses the semantic coherence of each duplicate group, allowing us to retain text-image pairs with the highest CLIP scores, as shown in Figure 3. This selective retention enhances the dataset by preserving the most relevant and semantically rich examples, minimizing redundancy while maintaining quality and diversity.

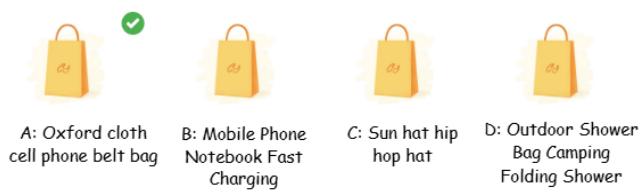


Figure 3: Quality Guided Deduplication retains the samples with better cross-modal alignment in duplicate groups to enhance the overall quality and achieve optimal data distributions.

Unimodal and Multimodal Curation Operators

EcoDatum enhances the quality of multimodal datasets by implementing rigorous unimodal and multimodal curation operators. The unimodal curation operators systematically filter out low-quality visuals and evaluate textual data for concreteness and relevance using both language identification and Image Caption Concreteness (ICC) metric (Yanuka et al. 2024). Multimodal curation integrates these approaches with advanced alignment techniques, employing models like GroundingDINO, an advanced open-set object detector (Liu et al. 2023) for precise local feature alignment and the CLIP model, for global semantic coherence. Together, these strategies ensure the curated dataset is of high quality, with well-aligned multimodal content.

Unimodal Curation Operators. For images, the specific heuristic operators filter out blurred and low-quality visuals. For texts, the FastText (Joulin et al. 2016) model identifies the language and the ICC metric evaluates the relevance and clarity of textual data using a pre-trained autoencoder.

Image-based quality filtering. Low-quality images can severely impact the learning of visual semantics. Our unimodal operators, based on heuristic rules, enhance dataset quality by filtering out images with detrimental attributes. The Geometric Operator targets images with non-standard aspect ratios that distort geometric relationships and compromise visual integrity when resized. Additionally, the DataComp dataset contains many intentionally blurred images to meet privacy standards, which reduces the visual detail crucial for effective model training. The Blurry Operator identifies and removes these excessively blurred images, ensuring that the curated dataset retains high visual quality.

Text-based caption assessment. We leverage the FastText model to identify and remove captions in rare languages, enhancing the linguistic consistency of our dataset. Additionally, we use the ICC metric, developed by a pre-trained autoencoder, to independently assess and filter captions. EcoDatum ensures the dataset retains only concrete and relevant captions, directly corresponding to their images.

Multimodal Curation Operators. EcoDatum enhances multimodal data curation by integrating both global and local image-text features, as shown in Figure 4. We employ GroundingDINO for precise local feature alignment, ensuring detailed correspondence between text and images at the object level. Additionally, we utilize the CLIP model, augmented with innovative adaptations, to maintain global semantic coherence throughout the dataset.

Local Cross-Modal Feature Alignment. We utilize GroundingDINO for the precise alignment of text descriptions with corresponding visual content. It integrates and analyzes text and visual data, effectively identifying relevant phrases in captions and accurately localizing associated visual elements within images, ensuring precise text-to-image alignment without prompt modification.

To quantitatively assess the alignment between text and images, we develop a metric based on the count of bounding

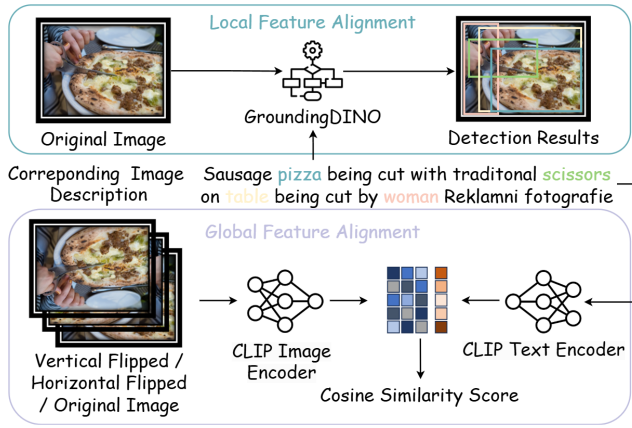


Figure 4: Illustration of multimodal curation operators integrating local and global cross-modal feature alignments.

boxes with confidence scores exceeding a predefined threshold, as shown in Eq (1). This metric serves to highlight the degree of correspondence between textual descriptions and visual representations. A higher count of accurate detections indicates richer, more detailed scenes, signifying that these data points are of higher value for training and subsequent applications. Data points that do not meet this threshold can be effectively filtered out, including those where the described objects do not visually correspond to the images or where the textual descriptions are insufficiently specific. This ensures our dataset excludes mismatches and generalities, retaining only high-quality, relevant multimodal content.

$$Count_{\text{GroundingDINO}} = \sum_{i=1}^n \{x_i > t\} \quad (1)$$

where x_i represents the confidence score of the i -th detected object, n represents the total number of objects detected, and t represents the predefined threshold.

This operator enhances the ability to curate multimodal data effectively, ensuring that the dataset maintains the most relevant and accurately aligned text-image pairs locally.

Global Cross-Modal Feature Alignment. In this module, EcoDatum utilizes the CLIP model, celebrated for its ability to assess the global semantic similarity between text descriptions and their visual counterparts. However, the effectiveness of the CLIP-Score can be compromised when images contain textual content that overlaps with captions. This issue is observed in 40% of the LAION dataset (Schuhmann et al. 2022) and 20% of the Datacomp dataset (Maini et al. 2023). To mitigate this, we implement an innovative adaptation known as Flip-CLIP, which includes Horizontal-CLIP (H-CLIP) and Vertical-CLIP (V-CLIP) techniques, inspired by (Yu et al. 2023). Before computing the CLIP scores, images are flipped horizontally or vertically, reducing the model’s bias towards text-based features and enabling more equitable evaluations of purely visual elements. The development of Flip-CLIP is motivated by the observation that OCR tasks often disproportionately influence the standard

CLIP score, especially when the image-text is overlapped.

By integrating both CLIP-Score and Flip-CLIP-Score, we foster the model’s ability to learn from visual content independently of textual influences, thereby enhancing EcoDatum’s capability to process and understand global visual features without excessive bias towards textual elements.

Modality Operators Ensemble

Given the vast volume of data and the high cost associated with obtaining high-quality labeled data, the availability of reliable labels is often limited. EcoDatum introduces a weak supervision labeling system that allows the efficient generation of quality-indicated labels at scale, mitigating the challenges of data scarcity and ensuring a more robust data quality assessment. In this study, data curation is abstracted as a data quality discrimination task, aiming to identify “high-quality” data. This ensemble-based system further enhances the capabilities of the data operators described above.

Specifically, EcoDatum employs a weak supervision ensemble model called LabelModel (Ratner et al. 2017; Bach et al. 2019) into the scope of data curation research, which integrates signal sources abstracted from unimodal and multimodal operators for data quality evaluation. This integration balances the limitations of individual operators and significantly reduces their erroneous impacts.

Each operator serves as an independent weak supervision signal source, assessing data quality from its unique dimension. The integration approach in this work uses LabelModel to combine multiple operators, automatically inferring a data quality score for each data sample by modeling the accuracy and relationships of these operators.

This process begins by matching each operator with corresponding labeling functions (LFs) (Ratner et al. 2017), which converts the operator’s inferred score s of the data sample x_i into weak supervision label L , as shown in Eq (2). The LFs compute operators’ inference results with the mean value b and the standard deviation β of the decision boundary to transform continuous scores into discrete labels. These labels are then aggregated to form a comprehensive weak supervision label matrix. In this context, weak supervision labeling with “Abstain” addresses situations where LFs face unclear features or inapplicable rules. Allowing the LabelModel to abstain from assigning labels in these cases prevents the generation of incorrect labels. This approach enhances the LabelModel’s ability to integrate diverse LFs by learning transformed matrix, particularly when they exhibit different biases and error patterns, thereby increasing the model’s robustness when handling heterogeneous data.

$$L_{x_i,j} = \begin{cases} 1, & \text{if } s_{x_i,j} \geq b_j + \beta_j \quad (\text{Selected}) \\ 0, & \text{if } s_{x_i,j} \leq b_j - \beta_j \quad (\text{Filtered}) \\ -1, & \text{if } b_j - \beta_j < s_{x_i,j} < b_j + \beta_j \quad (\text{Abstain}) \end{cases} \quad (2)$$

The LabelModel learns the transformed weak supervision label matrix L_M , estimating the weight w_j for each LF. These weights are used to combine the outputs of all LFs, ultimately generating a score for each data sample, which determines whether it is retained or filtered out. This

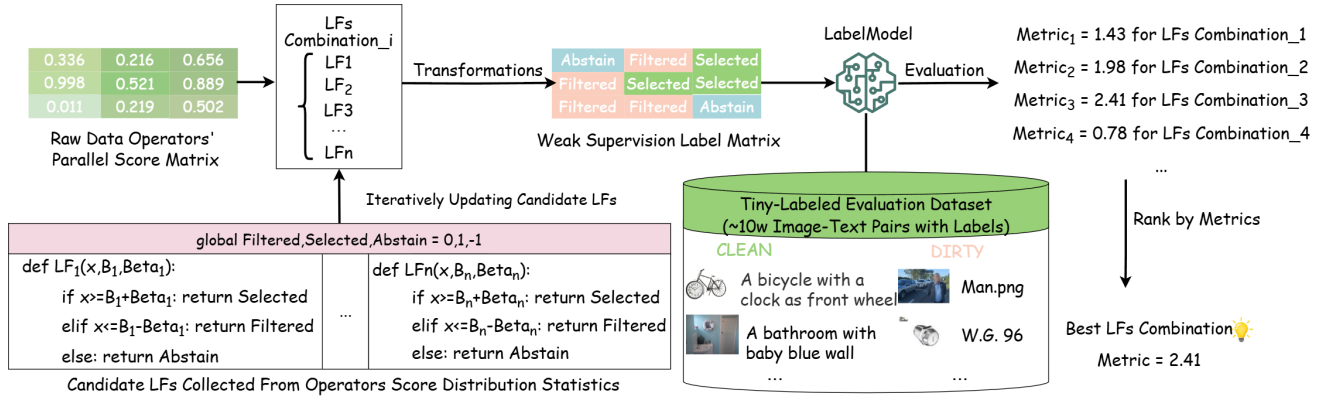


Figure 5: Overview of Search-Based LFs Combinations Optimization. This method optimizes LFs to create a more accurate weak supervision label matrix. The approach involves automatically searching optimal LFs to improve the matrix used by the LabelModel, evaluating LF combinations using a specialized composite metric over a gathered tiny-labeled dataset. This process boosts the LabelModel’s ability to assess data quality by refining operator interrelations and minimizing manual adjustments.

approach enhances the comprehensiveness and robustness of data quality evaluation, which ultimately allows the LabelModel to score all raw data and reflect quality.

Search-based Optimization

A novel search-based optimization method is introduced to enhance the design of LFs, improving the generation of a more accurate weak supervision label matrix for LabelModel modeling, as shown in Figure 5. This method addresses the challenge of converting operator-derived scores into labels by automatically optimizing LFs, reducing the need for manual experimentation. To further optimize the performance of the ensemble, EcoDatum proposes a composite metric that integrates the LabelModel’s data quality assessment capability with the attributes of LFs combination from the transformation steps, enabling a refined weak supervision label matrix. This approach enhances the LabelModel’s ability to analyze operator interrelations and importance, producing quality scores for data samples that closely approximate the ideal.

The evaluation stage automatically constructs a small labeled dataset containing “clean” and “noisy” samples. Clean data, labeled “1”, are sourced from the COCO dataset (Lin et al. 2014), while “noisy” samples, labeled “0”, are randomly sampled from the DataComp dataset to introduce both unimodal and multimodal noise and include added cross-modal noise through image-text pair exchanges. This setup tests the LabelModel’s ability to differentiate data quality via the F1-tiny scores in Eq (3). Importantly, this dataset is only used for assessing the LabelModel’s performance and does not contribute to training the model or optimizing Eq (3) coefficients, ensuring unbiased validation of the LF effectiveness.

To evaluate the data quality discrimination capacity of the LabelModel after learning generated weak supervision label matrices with different combinations of LFs, this research develops a specialized composite metric, shown in Eq (3), which combines classification metrics against ground

Algorithm 1: Search Optimization for LFs

- 1: **Input:**
 - 2: Raw dataset D_{raw} , Tiny-labeled dataset D_{tiny} , Operator Op , Evaluation metrics M , Candidate LF Combinations $LFsCombs$,
 - 3: **Output:**
 - 4: Optimal $LFsComb$
 - 5: Initialize M^*
 - 6: **for** each $LFs \in LFsCombs$ **do**
 - 7: Convert weak supervision label matrix L_M from LFs
 - 8: Train LabelModel on L_M
 - 9: Predict and evaluate LabelModel on D_{tiny}
 - 10: **if** $M(LFs) > M^*$ **then**
 - 11: Update $M^* \leftarrow M(LFs)$ and $LFs^* \leftarrow LFs$
 - 12: **end if**
 - 13: **end for**
 - 14: **return** LFs^*
-

truth and further incorporates the attributes of each operators’ LFs, specifically measuring the $f_{Overlap}$, $f_{Conflict}$, and $f_{Coverage}$. Here, they respectively indicate the frequency of agreement among LFs, the extent of disagreements, and the proportion of data labeled by at least one function.

$$M = \alpha_1 \cdot F1_{tiny} + \alpha_2 \cdot f_{Overlap} - \alpha_3 \cdot f_{Conflict} + \alpha_4 \cdot f_{Coverage} \quad (3)$$

Here, $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ are coefficients that are determined through a few rounds of experiments. These coefficients are tuned to optimize the balance between classification performance on the tiny labeled dataset and the contributions from overlap, conflict, and coverage metrics within the weak supervision labeling framework.

$$f_{Overlap} = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \left(\sum_{j=1}^m \mathbb{I}(LF_j(x_i) \neq 0) > 1 \right) \quad (4)$$

$$f_{\text{Conflict}} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(\exists j_1 \neq j_2, LF_{j_1}(x_i) \neq LF_{j_2}(x_i) \neq 0) \quad (5)$$

$$f_{\text{Coverage}} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}\left(\sum_{j=1}^m \mathbb{I}(LF_j(x_i) \neq 0) \geq 1\right) \quad (6)$$

These metrics capture the model’s effectiveness in integrating weak supervision signals. For example, even if the model achieves a high F1-score $F1_{\text{tiny}}$ on a tiny labeled dataset, significant conflicts between LFs or low coverage might still lead to instability in real-world applications. By optimizing these composite metrics, we can enhance the model’s generalization ability across different datasets.

The overall optimization process involves repeatedly constructing the LabelModel across candidate LFs combinations and using the results of the aforementioned composite metrics to identify the optimal LabelModel leaned by each weak supervision label matrix, as shown in Algorithm 1. This optimized model is then applied to the final data curation task. By employing the optimized LabelModel, the overall framework can maximize the robustness of data quality assessment, balance the limitations of individual operators, and ultimately enhance overall data efficiency.

Experiments

Dataset and Benchmark

The DataComp benchmark uniquely emphasizes data curation over model development. Unlike typical machine learning competitions that seek the best model with a fixed dataset, DataComp challenges participants to curate optimal datasets using fixed training code. This highlights the crucial role of high-quality, well-curated data in enhancing model performance. We choose the small-scale filtering track to validate the proposed framework EcoDatum, we curate a subset from a vast pool of 12.8 million image-text pairs from Common Crawl, adhering to the competition’s constraints of fixed training parameters and computational budgets. Our objective is to efficiently filter this dataset, ensuring consistency in training iterations regardless of dataset size.

The effectiveness of our curated dataset is evaluated across 38 diverse datasets, including ImageNet, 6 ImageNet distribution shift datasets, 12 datasets from the Visual Task Adaptation Benchmark, three retrieval datasets, and several others(Gadre et al. 2024). This extensive range of evaluation datasets tested the generalizability and robustness of EcoDatum, providing a comprehensive assessment of their impact on model training across various real-world scenarios.

Implementation Details

For the local cross-modal curation operator, we employ the GroundingDINO-based model(Liu et al. 2025) with Swin-Large as the image backbone and BERT-Base (Devlin 2018)

for encoding text, setting confidence thresholds at 0.1 to retain more potentially feature-aligned data. In global cross-modal curation, we use the CLIP-ViT-Large-14 architecture(Radford et al. 2021). In determining final data volume, we conducted extensive experiments and reviewed related works, concluding that approximately the top 40% samples by the EcoDatum quality score after deduplication (around 3.5M) provide the best balance between quality and quantity. Experiments utilized 8 NVIDIA A100 GPUs, The training and evaluation process required 2.5 hours. Data curation for the 12.8 million dataset involved approximately 10 hours.

Result Analysis

Existing Baselines. Several SOTA methods have previously set benchmarks in data filtering. LAION and CLIP Score utilize the CLIP model to refine datasets, while Datacomp Filtering employs heuristic unimodal operators for targeted data refinement (Gadre et al. 2023). The HYPeRbolic Entailment (HYPE) Filtering technique(Kim et al. 2024) enhances data quality by integrating unimodal specificity with cross-modal alignment. LINE’s strategy leverages large models for web data curation (Yokoo et al. 2023). The Text-Masking and Re-Scoring (T-MARS) method corrects imbalances where textual features overpower visual ones (Maini et al. 2023), and the University of Wisconsin-Madison’s (WS) approach utilizes an ensemble of object detection methods to optimize data filtering (Huang et al. 2024).

Performance Comparison. Building upon these foundations, EcoDatum enhances both efficiency and model training outcomes. As outlined in Table 1, using only 3.5 million data pairs from the original 12.8 million, EcoDatum achieved the highest average score of 0.182. This surpasses the performance of established methods like T-MARS and WS, both of which scored 0.180 across 38 diverse evaluation datasets. This curation strategy not only reduces computational overhead by 72% but also significantly improves data quality. EcoDatum exceeds the “No Filtering” baseline score of 0.132 and the Datacomp Basic filtering score of 0.142 by 28%. The integration of advanced methodologies like our optimized LabelModel for labeling functions tuning further refines the data curation process, setting new benchmarks in multimodal applications. The empirical results robustly validate our hypothesis that smaller, well-curated datasets can outperform larger, unfiltered datasets, underscoring the effectiveness of EcoDatum. Moreover, additional experiments show that EcoDatum consistently improves performance and scales effectively with increasing dataset size.

In this study, we introduce a composite metric designed to automatically optimize the generation of labeling functions (LFs), thereby facilitating the creation of a more accurate weak supervision label matrix. This optimization directly enhances the learning efficiency of the LabelModel, significantly improving its ability to assess data quality. To validate the effectiveness of this composite metric, we conducted a rigorous experimental case study. The process involved documenting a systematic search to identify the most effective LF combinations and repeatedly evaluating their impact on the average performance across a diverse set of 38 bench-

	No Filtering	LAION	Datacomp	CLIP	HYPE	LINE	T-MARS	WS	Ours
Dataset Size	12.8M	1.3M	3M	3.8M	2.3M	4.5M	2.3M	4.1M	3.5M
Avg. Perf.	0.132	0.133	0.142	0.173	0.176	0.177	0.180	0.180	0.182

Table 1: Performance comparison between our method, the Datacomp baseline, and other participants’ approaches.

mark tasks. The results, depicted in Figure 6, demonstrate a consistent positive correlation between the composite metric scores and the model’s performance, affirming the metric’s utility in refining the data curation process.

Ablation Study

This experiment conducts a systematic evaluation of data filtering techniques to assess impacts on the performance of the deep learning model, as detailed in Table 2. The “No Filtering” condition acts as the control group. “Random Deduplication” utilizes a stochastic method to eliminate duplicates, indicating that even indiscriminate reductions can improve model performance by balancing feature distribution.

Methods	Dataset Size	Avg. Perf.
No Filtering	12.8M	0.132
Random Dedup.	8.8M	0.145
Quality-Guided Dedup.(QGD)	8.8M	0.147
QGD+Ens.(Uni.)	3.5M	0.154
QGD+Ens.(Mul.)	3.5M	0.164
QGD+Ens.(Uni.&Global-Mul.)	3.5M	0.168
QGD+Ens.(Uni.&Local-Mul.)	3.5M	0.155
Best Perf.		
QGD+Ens.(Uni.&Mul.)	3.5M	0.182

Table 2: Performance comparison of different data curation and ensemble techniques over 38 downtasks.

The introduction of QGD achieves a 1.4% improvement over the random method with the same dataset size. Incorporating a unimodal operators’ ensemble within the QGD framework results in a 4.8% improvement, while a multimodal operators’ ensemble leads to a more substantial 9.5% enhancement. These results highlight the efficacy of both unimodal and multimodal operator ensembles in data curation. By integrating QGD with both unimodal and multimodal ensembles, the combined approach outperforms all others, showing a 45.4% improvement in performance compared to the “No Filtering” baseline. These experiments illustrate that EcoDatum strategically integrates advanced deduplication techniques and sophisticated ensemble frameworks to markedly elevate data quality, optimizing the pre-training process for multimodal models.

We conduct another ablation study to assess the individual contributions of data processing operators in data curation. By applying each operator independently and incrementally adding them, we explored their impact on downstream tasks. This approach allowed us to identify the most effective combinations of operators, significantly streamlining the optimization process. Through meticulous integration and refinement of labeling function (LF) constructions, we deter-

mined the most efficient operator combinations, thereby enhancing the accuracy and efficacy of our data curation methods. This conclusion suggests a strategic approach when dealing with massive web data and limited computational resources: focusing on alignment techniques can lead to more efficient data filtering. Such a focus can improve the generalization performance of multimodal models. Potentially, this experiment could pave the way for future research, indicating that more advanced image-text matching techniques might result in even better multimodal curation outcomes.

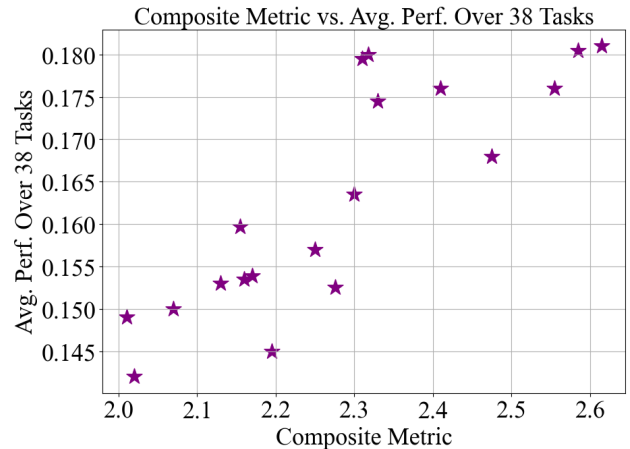


Figure 6: Composite Metric Validation with Repeated Experimental Downtasks Evaluations. The positive correlation indicates its capability to guide the tuning of the process.

Conclusion and Future Work

The volume of web-crawled datasets is rapidly expanding, and training multimodal models with such data are increasingly prevalent. This paper addresses the challenge of variable sample quality in web-crawled datasets by introducing a novel data curation framework, EcoDatum, designed to select high-quality data. EcoDatum begins with quality-guided deduplication to preprocess the data, followed by the integration of unimodal and multimodal operators into a weak supervision ensemble model, LabelModel, and have employed a search-based optimization method to refine the labeling matrix within LabelModel. Our experiments demonstrate robust performance across all evaluated tasks, securing a 1st place ranking in the small-scale track of the DataComp benchmark. While this study validated EcoDatum on a small dataset, future work will extend the evaluation to larger datasets. This expansion will further test the scalability of EcoDatum, aiming to solidify its effectiveness and efficiency in enhancing the training of multimodal models with diverse, large-scale web-crawled data.

Acknowledgments

This work was supported in part by China NSF grant No. 62202297 and in part by the Open Project Program of the Laboratory of Pinghu. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

References

- Abbas, A.; Tirumala, K.; Simig, D.; Ganguli, S.; and Morcos, A. S. 2023. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *arXiv preprint arXiv:2303.09540*.
- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Bach, S. H.; Rodriguez, D.; Liu, Y.; Luo, C.; Shao, H.; Xia, C.; Sen, S.; Ratner, A.; Hancock, B.; Alborzi, H.; et al. 2019. Snorkel drybell: A case study in deploying weak supervision at industrial scale. In *Proceedings of the 2019 International Conference on Management of Data*, 362–375.
- Beluch, W. H.; Genewein, T.; Nürnberger, A.; and Köhler, J. M. 2018. The power of ensembles for active learning in image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9368–9377.
- Breiman, L. 1996. Bagging predictors. *Machine learning*, 24: 123–140.
- Brown, T. B. 2020. Language models are few-shot learners. *arXiv preprint ArXiv:2005.14165*.
- Chen, D.; Huang, Y.; Ma, Z.; Chen, H.; Pan, X.; Ge, C.; Gao, D.; Xie, Y.; Liu, Z.; Gao, J.; et al. 2024. Data-juicer: A one-stop data processing system for large language models. In *Companion of the 2024 International Conference on Management of Data*, 120–134.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Devlin, J. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Farid, H. 2021. An overview of perceptual hashing. *Journal of Online Trust and Safety*, 1(1).
- Freund, Y.; and Schapire, R. E. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1): 119–139.
- Gadre, S. Y.; Ilharco, G.; Fang, A.; Hayase, J.; Smyrnis, G.; Nguyen, T.; Marten, R.; Wortsman, M.; Ghosh, D.; Zhang, J.; Orgad, E.; Entezari, R.; Daras, G.; Pratt, S. M.; Ramanujan, V.; Bitton, Y.; Marathe, K.; Mussmann, S.; Vencu, R.; Cherti, M.; Krishna, R.; Koh, P. W.; Saukh, O.; Ratner, A.; Song, S.; Hajishirzi, H.; Farhadi, A.; Beaumont, R.; Oh, S.; Dimakis, A.; Jitsev, J.; Carmon, Y.; Shankar, V.; and Schmidt, L. 2023. DataComp: In search of the next generation of multimodal datasets. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Gadre, S. Y.; Ilharco, G.; Fang, A.; Hayase, J.; Smyrnis, G.; Nguyen, T.; Marten, R.; Wortsman, M.; Ghosh, D.; Zhang, J.; et al. 2024. DataComp: In search of the next generation of multimodal datasets. *Advances in Neural Information Processing Systems*, 36.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; Casas, D. d. L.; Hendricks, L. A.; Welbl, J.; Clark, A.; et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Huang, T.-H.; Shin, C.; Tay, S. J.; Adila, D.; and Sala, F. 2024. Multimodal Data Curation via Object Detection and Filter Ensembles. *arXiv preprint arXiv:2401.12225*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, 4904–4916. PMLR.
- Joulin, A.; Grave, E.; Bojanowski, P.; Douze, M.; Jégou, H.; and Mikolov, T. 2016. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Kim, W.; Chun, S.; Kim, T.; Han, D.; and Yun, S. 2024. HYPE: Hyperbolic Entailment Filtering for Underspecified Images and Texts. *arXiv preprint arXiv:2404.17507*.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2025. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55. Springer.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Maini, P.; Goyal, S.; Lipton, Z. C.; Kolter, J. Z.; and Raghunathan, A. 2023. T-mars: Improving visual representations by circumventing text feature learning. *arXiv preprint arXiv:2307.03132*.
- Maini, P.; Yaghini, M.; and Papernot, N. 2021. Dataset inference: Ownership resolution in machine learning. *arXiv preprint arXiv:2104.10706*.
- Mindermann, S.; Brauner, J. M.; Razzak, M. T.; Sharma, M.; Kirsch, A.; Xu, W.; Höltingen, B.; Gomez, A. N.; Morisot, A.; Farquhar, S.; et al. 2022. Prioritized training on points that

- are learnable, worth learning, and not yet learnt. In *International Conference on Machine Learning*, 15630–15649. PMLR.
- Nguyen, T.; Gadre, S. Y.; Ilharco, G.; Oh, S.; and Schmidt, L. 2024. Improving multimodal datasets with image captioning. *Advances in Neural Information Processing Systems*, 36.
- Radenovic, F.; Dubey, A.; Kadian, A.; Mihaylov, T.; Vandenhende, S.; Patel, Y.; Wen, Y.; Ramanathan, V.; and Mahajan, D. 2023. Filtering, distillation, and hard negatives for vision-language pre-training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6967–6977.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rasmus, A.; Berglund, M.; Honkala, M.; Valpola, H.; and Raiko, T. 2015. Semi-supervised learning with ladder networks. *Advances in neural information processing systems*, 28.
- Ratner, A.; Bach, S. H.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB endowment. International conference on very large data bases*, volume 11, 269. NIH Public Access.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35: 25278–25294.
- Schuhmann, C.; Vencu, R.; Beaumont, R.; Kaczmarczyk, R.; Mullis, C.; Katta, A.; Coombes, T.; Jitsev, J.; and Komatsuzaki, A. 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2022. Learning from noisy labels with deep neural networks: A survey. *IEEE transactions on neural networks and learning systems*, 34(11): 8135–8153.
- Torkashvand, A.; Jameii, S. M.; and Reza, A. 2023. Deep learning-based collaborative filtering recommender systems: A comprehensive and systematic review. *Neural Computing and Applications*, 35(35): 24783–24827.
- Yanuka, M.; Alper, M.; Averbuch-Elor, H.; and Giryes, R. 2024. ICC: Quantifying Image Caption Concrete-ness for Multimodal Dataset Curation. *arXiv preprint arXiv:2403.01306*.
- Yokoo, S.; Zhu, P.; Ishikawa, Y.; Tanaka, M.; Kondo, M.; and Kataoka, H. 2023. Leveraging Image-Text Similarity and Caption Modification for the DataComp Challenge: Filtering Track and BYOD Track. *arXiv preprint arXiv:2310.14581*.
- Yu, H.; Tian, Y.; Kumar, S.; Yang, L.; and Wang, H. 2023. The devil is in the details: A deep dive into the rabbit hole of data filtering. *arXiv preprint arXiv:2309.15954*.
- Zimek, A.; Schubert, E.; and Kriegel, H.-P. 2012. A survey on unsupervised outlier detection in high-dimensional numerical data. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 5(5): 363–387.