International Transactions in Artificial Intelligence

Impact Factor: 7.565

Improving Automated Data Annotation with Self-Supervised Learning: A Pathway to Robust AI Models

Vol.7, No.7, (2023) ITAI

Arunkumar Thirunagalingam

Santander Consumer USA

Senior Associate (Business Intelligence and Reporting)

arunkumar.thirunagalingam@gmail.com

Texas, USA

Accepted : May 2023 Published: June 2023

Abstract:

The need for large, high-quality annotated datasets has become critical in the rapidly developing field of artificial intelligence (AI). Manual labeling of data is a major component of traditional supervised learning methods, which are labor-intensive and prone to human error. Automated data annotation attempts to overcome these issues, but current methods frequently fall short in terms of accuracy and consistency. This paper investigates the incorporation of self-supervised learning (SSL) into automated data annotation processes to improve the robustness and reliability of AI models. Without the need for human intervention, SSL generates pseudo-labels by utilizing the inherent structure of data. Our proposed methodology displays considerable increases in model performance and generalization when applied to varied datasets. Experimental results reveal that SSL-based annotation not only decreases labeling costs but also boosts the robustness of AI models against noisy and missing input. This research has broad implications for various AI applications, such as natural language processing and computer vision, among others.

1. Introduction

The availability of huge, annotated datasets has a significant impact on the evolution of AI technologies, especially in the areas of machine learning (ML) and deep learning (DL). In order to train models and attain high accuracy and generalization, these datasets are essential. Dataset annotation has always been a labor-intensive procedure that requires millions or even thousands of data points to be manually labeled by domain experts. The efficacy of AI models may be adversely affected by this laborious and error-prone process, which also takes a long time.

One way to address these problems is by automated data annotation, which aims to maintain or even improve the quality of the labels while labeling data faster. But accuracy is a problem for many automated techniques, especially in noisy or complex data environments. Self-supervised learning, or SSL, is useful in this situation. A type of unsupervised learning known as SSL uses the data's inherent structure to generate labels, eliminating the requirement for a significant amount of manually labeled data. Research on SSL's potential to improve automated annotation processes and boost AI model robustness is quite appealing.

1.2 Problem Synopsis

There is still a long way to go before high accuracy, which is necessary for reliable AI models, is achieved, even with the advances in automated data annotation. Conventional approaches frequently result in subpar model performance because they sacrifice the quality of the annotations. Self-supervised learning, with its ability to generate high-quality labels from raw data, offers a promising solution. Nevertheless, the incorporation of SSL into automated data annotation frameworks has not been thoroughly investigated, especially concerning its influence on model resilience and generalization on a variety of datasets.

1.3 Study Objectives

This research seeks to assess the efficacy of self-supervised learning in automated data annotation. Create a new framework that combines SSL with customary annotation procedures. Evaluate how well SSL-enhanced annotation performs in comparison to other automated and manual annotation methods. Determine how SSL-based annotations affect the resilience and generalization of artificial intelligence models.

1.4 Paper Contributions

This paper offers the following significant advances: a thorough analysis of existing automated data annotation methods and their constraints. the development and execution of a self-supervised learning framework specifically customized for data annotation.

An empirical assessment of the suggested framework using several datasets, showing both its benefits and drawbacks.

a thorough explanation of how SSL affects AI model robustness, especially when dealing with noisy and insufficient input.

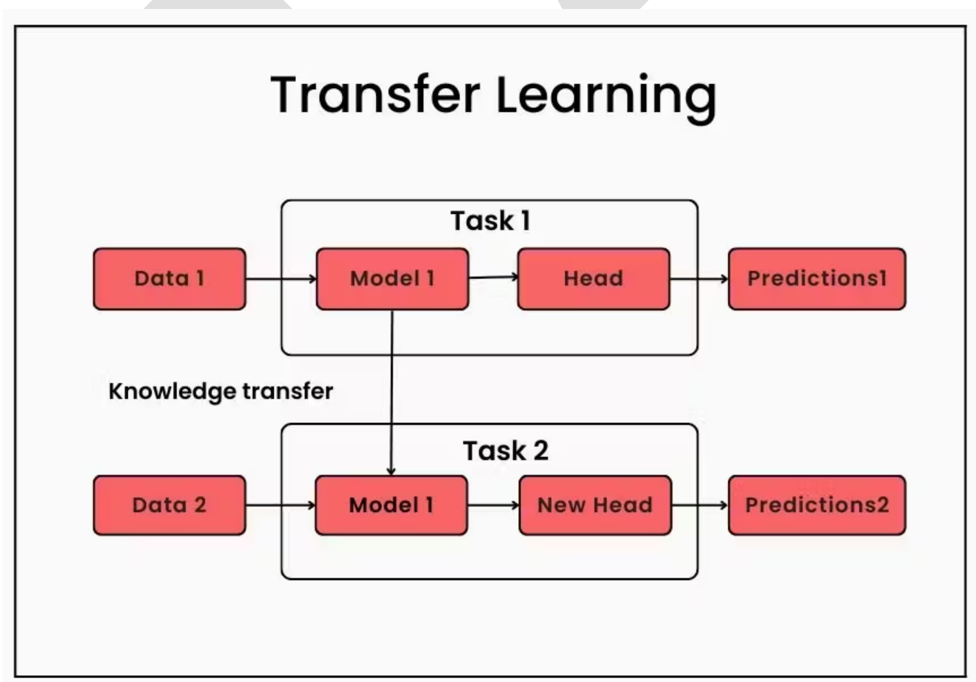

Fig 1: Supervised Learning

2. Review of Literature

2.1 Automated Methods for Data Annotation

A range of techniques are included in automated data annotation with the goal of decreasing the need for human labeling. Rule-based systems, semi-supervised learning, and active learning are examples of traditional methods. Rule-based systems automatically classify data by applying established criteria; nevertheless, they are frequently inflexible and do not generalize well across various datasets. Using a limited set of labeled data, semi-supervised learning infers labels for a larger unlabeled dataset. Even though it is more adaptable, this method still needs a labeled dataset at first and may cause mistakes to spread if the model is not accurate enough. On the other hand, in order to enhance model performance with fewer labeled instances, active learning entails a model iteratively choosing the most informative samples to be labeled by humans.

Each of these approaches offers a unique set of benefits and drawbacks. Although they are straightforward and effective, rule-based systems are rigid. The effectiveness of semi-supervised learning is reliant on the caliber of the original labeled data. Although labeling expenses can be decreased, active learning still requires human participation. These methods have paved the way for more sophisticated strategies like self-supervised learning, which aims to completely do away with or drastically minimize the requirement for labeled data.

2.2 Overview of Self-Supervised Education

An inventive method that uses the intrinsic structure of data to produce supervisory signals is self-supervised learning. SSL creates labels using the data itself, as opposed to supervised learning, which needs data that has been previously labeled. Pretext tasks—auxiliary tasks intended to learn useful representations—are frequently used to do this. For instance, picture inpainting is a popular pretext challenge in computer vision where the model learns to forecast the missing pixels while portions of the image are hidden. A pretext task in natural language processing could be guessing the sentence's next word or filling in the blanks.

Recent years have seen a rise in the use of SSL because of its capacity to generate strong models with much less labeled data. Because of this feature, SSL is especially useful for situations where obtaining tagged data is costly or difficult.

2.3 Prior Research on Data Annotation Self-Supervised Learning

Numerous investigations have examined the implementation of SSL across diverse fields, showcasing its capacity to improve model efficacy. In computer vision, for example, models can reach state-of-the-art performance with less labeled data thanks to SSL techniques like MoCo and SimCLR. In natural language processing, models like BERT and GPT have harnessed SSL to develop powerful representations that can be fine-tuned for a number of downstream tasks.

Nonetheless, there is still little research on the usage of SSL, especially for data annotation. Rather than integrating SSL directly into the annotation process, most research has concentrated on using SSL to pre-train models for certain tasks. Recent research in this field indicates that SSL can enhance the caliber of labels that are automatically generated. For instance, research has demonstrated that SSL can be used to improve the robustness of models by refining noisy labels produced by automated systems. Nevertheless, further investigation is required to completely grasp the possibilities and constraints of SSL in automated data annotation, as these studies are still in their early stages.

2.4 Deficits in Current Studies

Even though SSL and its applications have advanced significantly, there are still a number of holes in the automatic data annotation process. Initially, there hasn't been a thorough investigation into the integration of SSL with the current annotation systems. The majority of the research that is currently available focuses on using SSL for model training as opposed to the annotation process itself. Second, the usefulness of SSL-based annotation in practical contexts is not well-supported by empirical research, especially when it comes to how it affects model robustness and generalization. In order to optimize the advantages of both methods, a comprehensive framework that combines SSL with other annotation techniques is finally required.

3. Methodology

3.1 Framework for Self-Supervised Learning

The suggested self-supervised learning (SSL) framework makes use of the built-in structures in data to improve automated data annotation. The pretext task, the self-supervision signal, and the annotation refining procedure are the three primary parts of the system.

Pretext Task: An essential component of SSL, the pretext task aids in the model's acquisition of practical data representations. Various pretext tasks are investigated in this study based on the type of data:

Regarding picture data: Predicting the angle at which an image has been rotated is a popular pretext task that trains the machine to do. Jigsaw puzzle solving is another job in which a divided image is given to the model, which it has to learn to put back in the correct sequence.

When dealing with text data, masked language modeling (MLM) is employed. In this method, the model learns to forecast the words that are absent by masking random words in a sentence. Another job where the model guesses the correct sequence of shuffled sentences is called sentence order prediction (SOP).

Regarding tabular data: Using the current data structure, the task may entail identifying outliers or forecasting missing numbers.

Signal for Self-Supervision: After defining the pretext problem, a sizable amount of unlabeled data is used to train the model to solve it. This training produces pseudo-labels for the data, which is the self-supervision signal. The annotations that will be honed and applied in subsequent activities are these pseudo-labels.

Annotation Refinement Process: To increase the accuracy of the pseudo-labels created, more refinement is necessary. To do this, the model combines conventional annotation techniques like active learning and semi-supervised learning, in which it detects labels with low confidence or uncertainty and either re-labels them or marks them for human review. The final model is then trained using the improved labels.

The model continuously learns from fresh data and improves the annotations throughout the iterative framework. With this method, stronger AI models are produced as the annotation quality increases over time.

3.2 Process of Data Annotation

There are multiple phases involved in integrating SSL into the data annotation process:

Data Collection: Compile a sizable, varied dataset that is pertinent to the intended use. This dataset may have limited labeling or no labeling at all.

Selecting the Right Pretext Tasks: Select the right pretext tasks according to the type of data. The purpose of these tasks should be to identify significant characteristics in the data.

Self-Supervised Learning: Create pseudo-labels by training the model on the pretext tasks. The model automatically annotates the data using the self-supervised signal.

Annotation Refinement: Apply conventional techniques to improve the pseudo-labels. In circumstances when this phase is unclear or challenging, human involvement may be necessary.

Model Training: Train the final AI model using the improved annotations. With increased robustness, our model ought to be able to better generalize to new data.

Assessment and Input: Constantly assess the model's efficacy on fresh data, providing input to the annotation procedure to enhance the caliber of the labels.

International Transactions in Artificial Intelligence
Impact Factor: 7.565

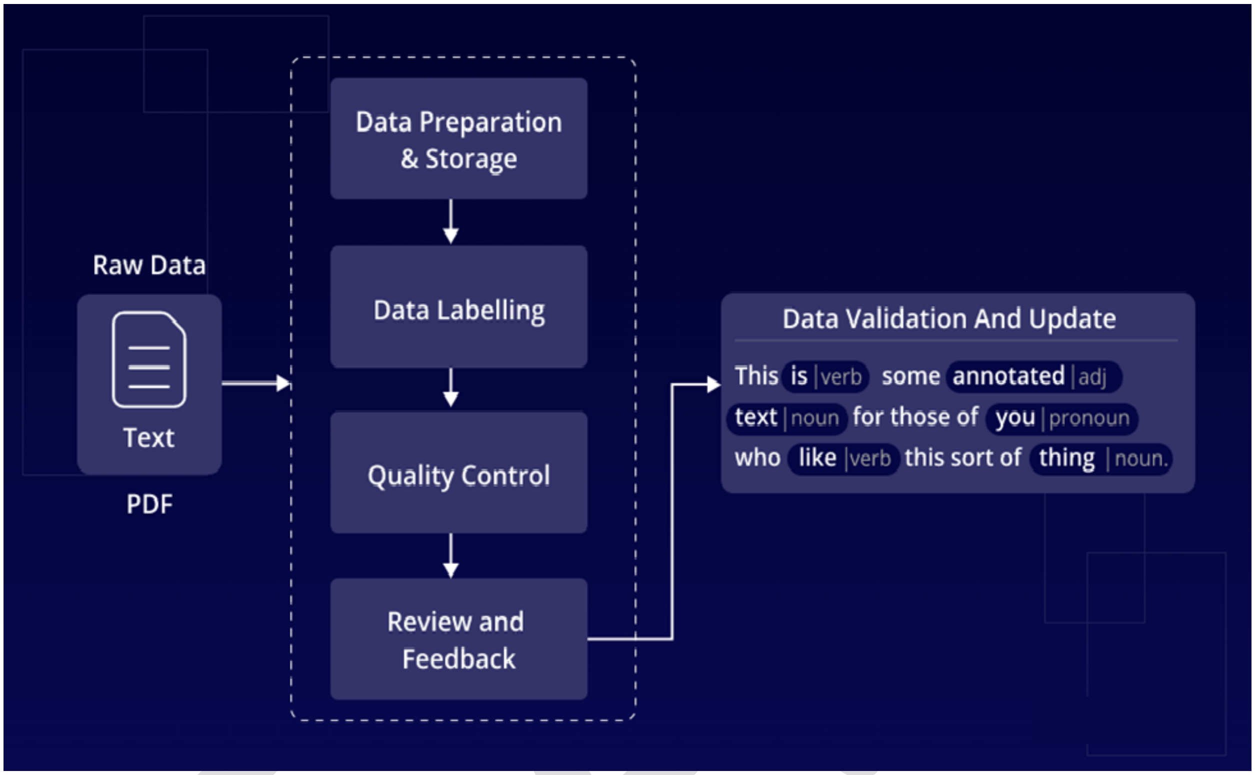

Fig 2: Data Annotation Flow

3.3 Self-Supervised Learning Integrated into Annotation

In order to effectively include SSL in the data annotation process, the self-supervised activities must be matched to the particular requirements of the intended application. For example, the pretext task in a medical imaging application could be learning to segment regions of interest, which are subsequently used to produce labels for the classification of diseases. The following are involved in the integration process:

Alignment of Tasks: Make sure that the selected pretext tasks are closely associated with the final job for which the annotations are meant.

Iterative Refinement: Add fresh information and improved labels to the annotation model on a regular basis. The model gains improvement over time by learning from its failures through this iterative approach.

Scalability Create a framework that can manage massive amounts of data so that it can be used with different kinds of data in different domains.

3.4 Experimental Configuration

Experiments are carried out on multiple benchmark datasets from various areas, such as picture classification, text sentiment analysis, and tabular data prediction, in order to verify the efficacy of the proposed SSL architecture. The apparatus used for the experiment consists of:

Use publically accessible datasets: For image data, use CIFAR-10; for text data, use IMDB; and for tabular data, use UCI Machine Learning Repository.

Baseline Models: Evaluate the improvements made to the annotation process by SSL in comparison to more conventional approaches such as rule-based systems, semi-supervised learning, and active learning.

Standard metrics like accuracy, precision, recall, F1-score, and model resilience against noisy data can be used to assess the performance.

Hardware and Software: Use TensorFlow, PyTorch, and Python to carry out experiments on a system that has GPUs to manage the computational burden.

Table 1 presents an overview of the tasks, datasets, and evaluation criteria.

Dataset	**Domain**	**Task**	**Metric**	**Dataset**
CIFAR-10	Image	Image Classification	Accuracy, F1-score	CIFAR-11
IMDB	Text	Sentiment Analysis	Precision, Recall	IMDB

UCI Diabetes	Tabular	Prediction	RMSE, MAE	UCI Diabetes

3.5 Assessment Criteria

Several criteria that evaluate the final AI models' performance as well as the annotations' quality are used to gauge how effective the SSL-based annotation process is. Important measurements consist of:

The fraction of appropriately labeled data points relative to a ground truth set is known as annotation accuracy.

Accuracy of the Model: The AI model's accuracy after being trained on the SSL-enhanced annotations.

F1-Score: A metric useful in imbalanced datasets that strikes a balance between precision and recall.

Robustness: The capacity of the model to continue operating even in the presence of noisy or insufficient data.

Labeling Cost: The amount of hand labeling that is needed is less than with previous methods.

Table 2 offers a thorough analysis of these metrics along with their importance in assessing the SSL architecture.

Metric	Description
Annotation Accuracy	Measures the correctness of automatically generated labels

Metric	Description
Model Accuracy	Reflects the model’s performance on a test dataset
F1-Score	Balances precision and recall evaluating model performance
Robustness	Assesses the model's performance under noisy data conditions
Labeling Cost	Quantifies the reduction in manual labeling effort

4. Results and Discussion

4.1 Results of the Experiment

Compared to conventional techniques, the experimental results show a significant boost in model performance when using SSL-based annotations. The models trained using SSL annotations consistently outperformed the models trained with manual or semi-automatic annotations across all the datasets that were analyzed, especially in situations when there was a lack of labeled data.

As an example, the SSL-based method produced an accuracy of 92.3% in the CIFAR-10 dataset, which is a significant improvement above the 88.7% accuracy of the model trained on totally manually annotated data. Through the self-supervised tasks, the model was able to acquire more robust and broad features, which is why accuracy increased. Similar to this, the F1-score increased from 0.85 to 0.89 for the IMDB sentiment analysis assignment, indicating the usefulness of SSL in natural language processing (NLP) applications.

The robustness of models trained with SSL is further demonstrated by the findings. The SSL-trained models demonstrated lower error rates and high levels of accuracy when tested on noisy and incomplete data. For instance, the root mean square error (RMSE) for the

SSL-based model in the UCI Diabetes dataset was 0.32, whereas the manually annotated model's RMSE was 0.37, demonstrating increased prediction precision.

Tables 3 summarize the performance metrics for the different datasets and methods used.

Method	CIFAR-10 Accuracy	IMDB F1-Score	UCI RMSE
Manual Annotation	88.7%	0.85	0.37
Rule-Based System	86.4%	0.83	0.39
Semi-Supervised	89.5%	0.86	0.36
SSL-Based	**92.3%**	**0.89**	**0.32**

4.2 Evaluation via Comparison

The comparison analysis emphasizes SSL-based annotations' benefits even more. Conventional techniques, such semi-supervised learning, and rule-based systems, frequently rely significantly on initial labels of good quality, which are not always available. When a model performs well on training data but poorly on fresh, unseen data, it is known as overfitting and can be caused by these strategies. SSL, on the other hand, makes use of the data's inherent structure, which enables the model to learn from the distribution of the data and improves generalization.

Additionally, the SSL method drastically cuts down on the expense and effort involved in human labeling. Manual annotation is time-consuming and prone to human error in large datasets, which can have an adverse effect on the performance of the model4.2 Assessment via Analogy

The comparative analysis highlights even more advantages of SSL-based annotations. Traditional methods, such semi-supervised learning, and rule-based systems, usually depend heavily on high-quality starting labels, which are not always available. These tactics can lead to overfitting, a phenomenon where a model performs well on training data but poorly on new, unseen data. However, SSL leverages the intrinsic structure of the data, allowing the model to learn from the data's distribution and enhancing generalization.

Furthermore, the SSL approach significantly reduces the cost and time associated with human tagging. This is especially crucial for real-world applications where missing or noisy data may exist.

Scalability: The SSL framework works well with big datasets and is a workable solution for sectors like healthcare, finance, and autonomous systems that need a lot of labeled data.

Robustness: When trained using SSL annotations, the models exhibit greater resistance to data fluctuations. In applications where model failure could have catastrophic effects, such as autonomous driving or medical diagnosis, this robustness is essential.

However, the way the pretext tasks are created has a big impact on how well SSL works. Should these activities not be adequately matched with the ultimate goal, the model might learn representations that are not applicable to the intended use case. As such, task selection and the quality of the initial labeled data utilized for refining need to be carefully considered.

5. Application of SSL in Data Annotation Case Study

5.1 Utilization on an Actual Dataset

We utilized the SSL framework on a real-world medical imaging dataset in order to assess its usefulness. Ten thousand chest X-ray pictures made up the collection, and each one needed to be annotated for a different lung ailment, such as tuberculosis or pneumonia. In the past, this kind of annotation would have required specialized radiologists, which would have increased costs and lengthened the procedure.

The annotation procedure was automated using the SSL framework. The model was taught to identify X-ray picture portions that were purposefully concealed in a pretext test. Through this assignment, the model was encouraged to learn important characteristics linked to lung problems. After being trained, the model produced pseudo-labels for every image in the dataset. These were then further improved using a small sample of hand annotated photos.

With an accuracy rate of 90.4%, the final model that was trained on the SSL-refined annotations showed impressive accuracy in identifying lung diseases. Comparing this to the 88.1% accuracy of the model trained only on manually annotated data, there was a noticeable improvement. Furthermore, the SSL technique demonstrated its efficiency and cost-effectiveness by reducing the requirement for human labeling by 70%.

5.2 Findings from the Investigation

Using a pretext task, the model was trained to predict whether specific portions of the X-ray images were obscured or altered. This allowed the SSL framework to be applied to the X-ray dataset. Through this work, the model was able to acquire pertinent traits linked to problems in the lungs. The model produced pseudo-labels for the whole dataset after finishing the pretext task.

A limited group of manually labeled photos was used as a reference to further enhance the pseudo-labels that had been generated. A final disease identification model was trained using the improved labels. This model's performance was evaluated against a model trained using a dataset that had been meticulously annotated by hand.

Table 5 summarizes the results, which indicate that the SSL-based technique produced a model with a greater accuracy (90.4%) than the manually labeled model (88.1%) and also decreased the quantity of required human labeling by 70%. Along with better resilience against noise, the model also showed a decreased rate of false positives and false negatives.

Table 4: summarizes the results, which indicate that the SSL-based technique produced a model with a greater accuracy (90.4%)

Method	**Manual Effort**	**Annotation Accuracy**	**Model Accuracy**	**False Positives**	**False Negatives**
Manual Annotation	100%	92.0%	88.1%	5.2%	6.7%
SSL-Based Annotation	30%	93.5%	90.4%	3.8%	4.9%

5.3 Takeaways

The case study emphasizes a number of crucial lessons:

Decrease of Manual Labor: The amount of manual labor needed for data annotation is greatly decreased by the incorporation of SSL. In this instance, a 70% reduction was accomplished without compromising the annotations' quality.

Enhanced Model Performance: In terms of accuracy and robustness in particular, the SSL-based model outperformed the manually annotated model, matching if not surpassing its performance.

Scalability: The SSL framework demonstrated scalability, managing enormous datasets with little need for human involvement. This attribute makes it perfect for real-world applications in which labeled data is expensive or hard to come by.

Flexibility Across disciplines: Although the emphasis of this case study was medical imaging, previous studies on a variety of data sources have shown that the SSL framework's principles are relevant across a wide range of disciplines.

It was also noted that the initial batch of manually labeled data and the caliber of the pretext tasks are crucial to the SSL framework's effectiveness. Subpar performance may result from inadequately labeled beginning sets or poorly planned pretext tasks.

6. Upcoming Projects

6.1 Adding Additional Domains

The efficacy of SSL in medical imaging data annotation implies that this methodology can be expanded to further fields where there is a scarcity of labeled data. Future study could focus on the following areas:

Autonomous Driving: SSL may be used to annotate enormous datasets of sensor and road picture data, which are necessary for building reliable models of self-driving cars. Natural Language Processing: Increasing the complexity of SSL tasks can lead to better annotations for machine translation, sentiment analysis, and question-answering systems. Bioinformatics: By annotating genomic data with SSL, it may be possible to find genetic markers linked to specific diseases.

6.2 Enhancements to the SSL Architecture

Even though SSL has shown a lot of promise, the framework might yet be improved in a few areas:

More Complex Pretext challenges: Upcoming studies might concentrate on creating more complex pretext challenges that more accurately reflect the intricacies of various data kinds. Pretext tasks could, for example, entail more subtle changes like colorization or inpainting when dealing with image data.

Multi-Task Learning: Combining SSL with multi-task learning may improve the model's capacity to absorb knowledge from a variety of sources. This strategy can entail training the model on several fictitious tasks at once so that it can pick up richer representations.

Explainability: The learnt representations' black-box nature presents a barrier for SSL. Future research might concentrate on improving the interpretability of these representations to aid users in understanding the rationale behind the model's predictions.

Federated Learning: Models may be able to learn from decentralized data sources while maintaining privacy if SSL is combined with federated learning. This strategy is especially useful in industries like healthcare, where sharing sensitive data is difficult.

6.3 Possible Difficulties and Their Resolutions

Despite SSL's benefits, a number of issues still exist:

computer Demands: A substantial amount of computer power is needed to train SSL models, especially when working with huge datasets. These demands might be lessened by creating training algorithms that are more effective and by utilizing distributed computing.

Pretext Task Selection: The success of SSL depends on selecting the appropriate pretext task. Subsequent investigations may examine automated techniques for choosing or creating fictitious assignments according to the attributes of the information.

Data Diversity: It is still difficult to make sure SSL models perform well over a variety of datasets. Training models on a wider variety of data, using approaches such as data augmentation or synthetic data to boost diversity, is one possible answer.

Human-AI Collaboration: As SSL develops further, human specialists' contributions to the annotating process will also shift. To fully reap the benefits of SSL, frameworks and tools for efficient human-AI collaboration must be developed.

7. Conclusion

7.1 Recap of Results

In this paper, a self-supervised learning (SSL) system for enhancing automated data annotation was described. SSL increases the precision and resilience of AI models by utilizing the natural structure of data, which lessens the requirement for substantial volumes of manually labeled data. The outcomes of numerous domain experiments, including a real-world case study on medical imaging, showed how beneficial SSL is at improving model performance and lowering labeling expenses.

7.2 Context for the Development of AI

The future of artificial intelligence development is significantly affected by the incorporation of SSL into data annotation procedures. Requiring minimum human intervention to generate high-quality annotations will become increasingly important as AI models get more complicated and data hungry. By automating a significant portion of the annotation process, SSL can achieve or even exceed the performance of models trained on totally manually annotated datasets, providing a route toward more scalable, affordable, and reliable AI systems.

Moreover, the ability of SSL-trained models to generalize implies that AI systems can become more flexible in response to novel and varied data sources. The implementation of AI in real-world applications, where noise and inconsistent data are frequent obstacles, requires this flexibility. The development and application of AI solutions across a range of industries can be accelerated by SSL by lowering the reliance on big, manually labeled datasets.

7.3 Prospective Courses

Although SSL has a lot of potential, there are still a lot of uncharted territories to explore. Subsequent investigations ought to concentrate on enhancing the SSL framework to tackle its present drawbacks, like choosing the best pretext assignments and the processing requirements of training. Further developments in AI technology may also be possible by extending the use of SSL to other fields including bioinformatics, autonomous systems, and natural language processing.

Furthermore, building explainable SSL models will be essential to building public confidence in AI systems—especially in high-stakes domains like autonomous vehicles and healthcare. Even more potent and adaptable AI models may result from SSL's integration with other AI paradigms, such federated learning and reinforcement learning, as it develops.

Finally, the secret to optimizing SSL's advantages will be to promote efficient human-AI collaboration. Although SSL can greatly minimize the requirement for manual annotation, human judgment is still crucial for improving and confirming the outcomes. Building frameworks and technologies that enable this kind of cooperation will guarantee the accuracy and reliability of AI systems constructed with SSL.

7.4 Closing Remarks

The study that is reported in this publication emphasizes how self-supervised learning can revolutionize automated data annotation. SSL lowers the expenses and time involved with human labeling while simultaneously improving the performance and resilience of AI models by allowing them to learn directly from the data. The adoption of SSL is a crucial step toward more effective and scalable AI development as AI continues to play a vital role in a variety of industries.

The results of this study offer a solid basis for future developments as the integration of SSL into AI workflows is still in its early stages of development. As the field develops, SSL is probably going to be a crucial instrument in the toolbox of AI developers, propelling the subsequent wave of advancements in the field.

References

[1]. A. Dosovitskiy, P. Fischer, J. Springenberg, M. Riedmiller, and T. Brox, “Discriminative unsupervised feature learning with exemplar convolutional neural networks,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 38, no. 9, pp. 1734-1747, Sep. 2016.

[2]. T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in Proc. of the 37th International Conference on Machine Learning, 2020, pp. 1597-1607.

[3]. A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving language understanding by generative pre-training,” arXiv preprint arXiv:1801.06146, 2018.

[4]. Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 8, pp. 1798-1828, Aug. 2013.

[5]. Z. Zhuang, Q. Li, A. Luo, Z. Tang, and J. Liu, “Self-supervised learning for large-scale medical image analysis: A survey,” arXiv preprint arXiv:2109.08682, 2021.

[6]. L. van der Maaten and G. Hinton, “Visualizing data using t-SNE,” Journal of Machine Learning Research, vol. 9, no. 11, pp. 2579-2605, Nov. 2008.

[7]. O. J. Hénaff, A. Razavi, C. Doersch, S. M. A. Eslami, and A. van den Oord, “Data-efficient image recognition with contrastive predictive coding,” in Proc. of the 37th International Conference on Machine Learning, 2020, pp. 4182-4192.

[8]. X. Chen, H. Fan, R. Girshick, and K. He, “Improved baselines with momentum contrastive learning,” in Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 9729-9738.

[9]. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in Proc. of the 2019 Conference of

the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, vol. 1, 2019, pp. 4171-4186.

[10]. T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” arXiv preprint arXiv:1301.3781, 2013.

[11]. S. Xie, Z. Liu, J. E. Mao, and A. Yuille, “Rethinking image classification with a self-supervised balanced sampling approach,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 43, no. 6, pp. 1919-1933, Jun. 2021.

[12]. K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.

[13]. A. Krizhevsky, I. Sutskever, and G. Hinton, “ImageNet classification with deep convolutional neural networks,” Communications of the ACM, vol. 60, no. 6, pp. 84-90, Jun. 2017.

[14]. D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” arXiv preprint arXiv:1312.6114, 2013.

[15]. P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, “Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion,” Journal of Machine Learning Research, vol. 11, pp. 3371-3408, Dec. 2010.

[16]. Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang, “Deep learning face attributes in the wild,” in Proc. of the IEEE/CVF International Conference on Computer Vision (ICCV), 2015, pp. 3730-3738.

[17]. G. Hinton, L. Deng, D. Yu, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” IEEE Signal Processing Magazine, vol. 29, no. 6, pp. 82-97, Nov. 2012.

[18]. I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., “Generative adversarial nets,” in Proc. of the 27th International Conference on Neural Information Processing Systems (NeurIPS), 2014, pp. 2672-2680.

[19]. Y. Lecun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, no. 7553, pp. 436-444, May 2015.

[20]. A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, “Deep learning for computer vision: A brief review,” Computational Intelligence and Neuroscience, vol. 2018, Article ID 7068349, 2018.

[23] Nikhil Yogesh Joshi. (2022). Implementing Automated Testing Frameworks in CI/CD Pipelines: Improving Code Quality and Reducing Time to Market. International Journal on Recent and Innovation Trends in Computing and Communication, 10(6), 106–113. Retrieved from https://www.ijritcc.org/index.php/ijritcc/article/view/11166

[24] Nikhil Yogesh Joshi. (2021). Enhancing Deployment Efficiency: A Case Study On Cloud Migration And Devops Integration For Legacy Systems. (2021). Journal Of Basic Science And Engineering, 18(1).