



A Survey and Approach to Chart Classification

Anurag Dhote^{1,3}, Mohammed Javed^{1,3}, and David S. Doermann²

¹ Indian Institute of Information Technology Allahabad, Prayagraj, India
{mit2021082,javed}@iiita.ac.in

² Department of CSE, University at Buffalo, Buffalo, NY, USA
doermann@buffalo.edu

³ Computer Vision & Biometrics Lab, Dept. of IT, IIIT Allahabad, Prayagraj, India

Abstract. Charts represent an essential source of visual information in documents and facilitate a deep understanding and interpretation of information typically conveyed numerically. In the scientific literature, there are many charts, each with its stylistic differences. Recently the document understanding community has begun to address the problem of automatic chart understanding, which begins with chart classification. In this paper, we present a survey of the current state-of-the-art techniques for chart classification and discuss the available datasets and their supported chart types. We broadly classify these contributions as traditional approaches based on ML, CNN, and Transformers.

Furthermore, we carry out an extensive comparative performance analysis of CNN-based and transformer-based approaches on the recently published CHARTINFO UB-UNITECH PMC dataset for the CHART-Infographics competition at ICPR 2022. The data set includes 15 different chart categories, including 22,923 training images and 13,260 test images. We have implemented a vision-based transformer model that produces state-of-the-art results in chart classification.

Keywords: Chart Classification · Deep Learning · Chart Mining

1 Introduction

Charts provide a compact summary of important information or research findings in technical documents and are a powerful visualization tool widely used by the scientific and business communities. In the recent literature, the problem of chart mining has attracted increased attention due to numerous advantages, as suggested in the comprehensive survey published by Davila et al. in 2019 [11]. The term Chart mining refers to the process of extracting information represented by charts. Another motivating factor in the increased attention paid to this problem is a series of competitions held in conjunction with significant conferences to address the critical challenges in the chart mining pipeline [10, 12, 13].

Since a variety of charts are possible, chart classification is often the first step in chart mining. The task of chart image classification can be formalized as, given

a chart image extracted from a document, classifying the image into one of N defined categories. The wide variety of chart types in the literature adds to the complexity of the task [6, 11, 34]. Some additional problems include interclass similarity, noise in authentic chart images, and more state-of-the-art datasets that cover multiple chart types and incorporate 2.5 or 3D charts and noise into the training samples [34]. The rise of robust deep learning models has contributed significantly to the success of chart classification. Deep learning approaches have outperformed traditional machine learning approaches regarding robustness and performance. Yet there need to be more state-of-the-art solutions that can provide stable results and are robust enough to address noise in some data sets. In this paper, we provide a performance comparison of several deep learning models that are state-of-the-art in the ImageNet [28] classification task. In addition, we report the performances of several popular vision transformers, which, to the best of our knowledge, have yet to be used for chart classification, except for the recent ICPR 2022 CHART-Infographics competition [13].

This paper is organized as follows. Section 2 summarizes the existing chart classification literature covering traditional and deep learning-based methods, including a brief discussion on transformer-based chart classification. Section 3 reports and summarizes publicly available datasets. Section 4 briefly highlights the popular ImageNet pre-trained deep learning-based models that will be used for our comparative study. Section 5 describes the latest edition of the UB PMC dataset, the training and testing protocols, and a discussion on their performance for chart classification. Section 6 provides information on possible improvements and suggestions for future research. Finally, Sect. 7 concludes with a summary of the paper.

2 Chart Classification Techniques

Based on the type of approaches used to implement the chart classification task in the literature, they can be grouped into traditional ML, CNN-based deep learning, and Transformer-based deep learning. Each type of approach is described briefly below.

2.1 Traditional ML Approaches

Traditional approaches rely on feature extraction methods that are often manual and general-purpose. Features are extracted and then represented in mathematical form for direct processing by machine learning classifiers. Savva et al. [29] present a system that automatically reformats visualizations to increase visual comprehension. The authors use low-level image features for classification in conjunction with text-level features. The system uses a multiclass SVM classifier trained on a corpus containing 2601 chart images labeled with ten categories, following Gao et al.’s manual extraction approach. In [14], researchers propose VIEW, a system that automatically extracts information from raster-format

charts. The authors used an SVM to separate the textual and graphical components and classify the chart images based on the graphic elements extracted from the visual components. The text is typically found in three chart categories - bar charts, pie charts, and line graphs, with 100 images for each category collected from various real-world digital resources.

Instead of taking an image as input, Karthikeyani and Nagarajan [19] present a system to recognize chart images from PDF documents using eleven texture features that are part of a Gray Level Co-Occurrence Matrix. A chart image is located in the PDF Document database, and the features are extracted and fed to the learning model. SVM, KNN, and MLP are the classifiers used for classification. Cheng et al. [7] employ a multimodal approach that uses text and image features. These features are provided as input to an MLP. The output is characterized as a fuzzy set to get the final result. The corpus contains 1707 charts with three categories and a 96.1% classification result.

2.2 CNN-Based Deep Learning Approaches

Liu et al. [22] used a combination of Convolutional Neural Networks (CNNs) and Deep Belief networks (DBNs) to capture high-level information present in deep hidden layers. Fully Connected Layers of Deep CNN are used to extract deeply hidden features. A DBN is then used to predict the image class using the deep hidden features. The authors use transfer learning and perform fine-tuning to prevent overfitting. They use a data set that includes more than 5,000 images of charts, including pie, scatter, line, bar, and flow classes. Deep features are useful over primitive features to provide better stability and scalability to the proposed framework. The proposed method achieves an average accuracy of 75.4%, which is 2.8% more than the method that uses only deep ConvNets.

Given the results of CNN in the classification of natural images, Siegel et al. [30] used two CNN-based architectures for chart classification. They evaluated AlexNet and ResNet-50, which are pre-trained on the ImageNet data set and then fine-tuned for chart classification. This transfer learning approach is prevalent in subsequent works addressing this particular problem. The proposed frameworks outperformed the state-of-the-art model at the time, such as ReVision, by a significant margin. ResNet-50 achieved the best classification accuracy of 86% on a data set that contained more than 60000 images spread over seven categories.

Amara et al. [1] proposed a CNN-based on LeNet to classify images from their corpus of 3377 images into 11 categories. The model comprises eight layers, one input layer, five hidden layers, one fully connected layer, and one output layer. The fully connected layer is used as a classifier, while the hidden layers are convolution and pooling layers designed to extract features automatically. A fully connected layer employs softmax activation to classify images into defined classes. For evaluation of the model's performance, an 80-20 split is performed on the data set for training and assessment. The proposed model performs better than the LeNet and pretrained LeNet architectures with an accuracy of 89.5%.

Jung et al. [18] present a classification method using the deep learning framework Caffe and evaluate its efficacy by comparing it with ReVision [29]. The authors use GoogLeNet [32] for classification and compare its results with shallower networks like LeNet-1 and AlexNet [20]. GoogLeNet outperforms LeNet-1 and AlexNet with an accuracy of 91.3%. Five-fold cross-validation is used for calculating the accuracy on an image corpus with 737–901 images for each chart type. The test concludes that ChartSense provides higher classification accuracy for all chart types than ReVision.

With studies adapting the deep learning approach for chart image classification, a comparative study of traditional vs. CNN architectures was required. Chagas et al. [6] provide a comparative analysis of conventional vs. CNN techniques. Authors evaluated CNN architectures (VGG19 [31], Resnet-50 [15], and Inception-V3 [33]) for chart image classification for ten classes of charts. The performance is compared with conventional machine learning classifiers, Naive Bayes, HOG features combined with KNN, Support Vector Machines, and Random Forests. Pre-trained CNN models with fine-tuned last convolutional layers were used. The authors concluded that CNN models surpass traditional methods with an accuracy of 77.76% (Resnet-50) and 76.77% (Inception-V3) compared to 45.03% (HOG + SVM).

Dia et al. [9] employ four deep learning models on a corpus of 11,174 chart images of five categories. Of AlexNet [20], VGG16 [31], GoogLeNet [32] and ResNet [15], the authors get the best accuracy of 99.55% for VGG16 model. VGG16 outperforms the models used in ChartSense paper by a large margin.

Significant roadblocks to chart mining research are caused by the fact that current chart data sets must be larger and contain sufficient diversity to support deep learning. To address this problem, Jobin et al. [21] presented DocFigure, a chart classification data set with 33,000 charts in 28 different classes. To classify charts, the author’s proposed techniques utilize deep features, deep texture features, and a combination of both. Among these baseline classification techniques, the authors observed that combining deep features and deep texture features classifies images more efficiently than individual features. The average classification accuracy improved by 3.94% and 2.10% by concatenating FC-CNN and FV-CNN over individual use of FC-CNN and FV-CNN, respectively. The overall accuracy of the combined feature methods turned out to be 92.90%.

Luo et al. proposed a unified method to handle various chart styles [26], where they show that generalization can be obtained in deep learning frameworks with rule-based methods. The experiments were performed on three different datasets of over 300,000 images with three chart categories. In addition to the framework, an evaluation metric for the bar, line, and pie charts is also introduced. The authors concluded that the proposed framework performs better than traditional rules-based and pure deep learning methods.

Araújo et al. [2] implemented four classic CNN models that performed well on computer vision tasks, including Xception [8], VGG19 [31], ResNet152 [15] and MobileNet [16]. The weights of these models were pre-trained on the ImageNet dataset, and the authors further performed hyperparameter tuning to obtain a stable learning rate and weight decay. These models were employed on

a self-aggregated chart image corpus of 21,099 images with 13 different chart categories. Xception outperforms the other models by hitting an accuracy of 95%.

The problem of small datasets has been prevalent since the problem of chart mining was first introduced. Most work tries to increase the size of the dataset. However, Bajic and Job [4] use a Siamese CNN network to work with smaller datasets. The authors show that an accuracy of 100% can be achieved with 50 images per class, which is significantly better than using a vanilla CNN.

With the increase in datasets for chart images and the rise of deep learning models being employed on said datasets, an empirical study of these deep learning models was due. Thiyam et al. [35] compared 15 different deep-learning models on a self-aggregated dataset of 110,182 images spfeatures24 different chart categories. In addition, the authors tested the performance of these models on several preexisting test sets. They concluded that Xception (90.25%) and DenseNet121(90.12%) provide the most consistent and stable performance of all the deep learning models. The authors arrived at this decision by employing a five-fold cross-validation technique and calculating the standard deviation for each model across all datasets.

Davila et al. [10] summarized the work of different participants in the competition’s first edition by harvesting raw tables from Infographics that provided data and tools for the chart recognition community. Two data sets were provided for the classification task. One was a synthetically generated AdobeSynth dataset, and the other UB-PMC data set was gathered from the PubMedCentral open-access library. The highest average F1-measure achieved for the synthetic data set was 99.81% and the highest F1-measure achieved for the PMC data set was 88.29%. In the second edition of the competition, the PMC set was improved and included in the training phase. An ensemble of ResNet152 and DenseNet121 achieved the highest F1-score of 92.8%. The third edition of the competition was recently held at ICPR 2022. The corpus of real chart images was made up of 36,183 chart images. The winning team achieved an F1 score of 91% with a base Swin transformer model with a progressive resizing technique. We summarize the competition details in Table 1.

Table 1. Competition on Harvesting Raw Tables from Infographics (CHART-Infographics)

Competition	Dataset	#Classes	Train Size	Test Size	Top performing Model	F1-measure
ICDAR 2019 [10]	AdobeSynth	10	198,010	4540	ResNet-101	99.81%
	PMC	7		4242		88.29%
ICPR 2020 [12]	Adobe Synth	12	14,400	2,999	DenseNet-121	100%
	UB PMC	15			15,636	
ICPR 2022 [13]	UB PMC	15	22,923	13,620	Swin Transformer	91%

Table 2. Published Literature on Chart Classification

Authors	Dataset	Model	Metric	Performance
Savva et al. [29]	Self-acquired	SVM	Accuracy	96.00%
Gao et al. [14]	Self-acquired	SVM	Accuracy	97.00%
Kartikeyani and Nagarajan [19]	Self-acquired	MLP	Accuracy	69.68%
		KNN		78.06%
		SVM		76.77%
Cheng et al. [7]	Self-acquired	MLP	Accuracy	96.10%
Liu et al. [22]	DeepChart	CNN + DBN	Accuracy	75.40%
Siegel et al. [30]	ChartSeer	AlexNet	Accuracy	84.00%
		ResNet-50		86.00%
Amara et al. [1]	Self-acquired	CNN	Accuracy	89.50%
Jung et al. [18]	Chart-Sense	GoogLeNet	Accuracy	91.30%
Balaji et al. [5]	Self-acquired	CNN	Accuracy	99.72%
Chagas et al. [6]	Chart-Vega	ResNet-50	Accuracy	76.76%
		Inception-V3		76.77%
Dai et al. [9]	Self-acquired	ResNet	Accuracy	98.89%
		GoogLeNet		99.07%
		AlexNet		99.48%
		VGG-16		99.55%
Liu et al. [23]	Self-acquired	VGG-16	Accuracy	96.35%
Davila et al. [10]	Synthetic	ResNet-101	F1-measure	99.81%
	UB-PMC	ResNet-101		88.29%
Jobin et al. [21]	DocFigure	FC-CNN + FV-CNN	Accuracy	91.30%
Bajic et al. [3]	Self-acquired	VGG-16	Accuracy	89.00%
Araujo et al. [2]	Self-acquired	Xception	Accuracy	95.00%
Luo et al. [26]	Chart-OCR	CNN	Custom(Bar)	91.90%
			Custom(Pie)	91.80%
			Custom(Line)	96.20%
Davila et al. [12]	UB-PMC	DenseNet-121 + ResNet-152	F1-measure	92.80%
Bajic and Job [4]	Self-acquired	Siamese CNN	Accuracy	100%
Thiyam et al. [35]	Self-acquired	Xception	Accuracy	90.25%
		DenseNet121		90.12%
		DenseNet201		90.53%
Davila et al. [13]	UB-PMC	Swin Transformer	F1-measure	91.00%

2.3 Transformer-Based Deep Learning Approaches

Since the inception of Vision Transformer, there has been a lot of development in various computer vision tasks such as image classification, object detection, and image segmentation. Vision transformer has outperformed CNN-based models in these tasks on the ImageNet dataset. However, there has not been widespread application of vision transformers to chart image classification. To our knowledge, only the Swin transformer [24] has been used for chart classification as reported in [13], which won the CHART-Infographics challenge ICPR2022. The authors applied a Swin Transformer Base Model with a progressive resizing technique. The models were initially trained on a scale (input size) of 224 followed by 384 [13].

The existing models in the literature are summarised in Table 2.

3 Chart Classification Datasets

There has been a significant increase in the size of datasets both in terms of the number of samples and the number of chart types. The Revision dataset [29] had only 2,601 images and 10 chart types. The recent publicly available dataset [13] comprises around 33,000 chart images of 15 different categories. The details of several publicly available datasets are discussed in this section.

Table 3. Chart Classification Datasets

Dataset	Year	#Samples	#Category	Public (Y/N)
ReVision [29]	2011	2601	10	Y
View [14]	2012	300	3	N
Self [19]	2012	155	8	N
Self [7]	2014	1707	3	N
DeepChart [22]	2015	5000	5	Y
ChartSeer [30]	2016	60000	7	N
Self [1]	2017	3377	11	N
Chart-Sense [18]	2017	6997	10	Y
Chart-Text [5]	2018	6000	2	N
Chart-Vega [6]	2018	14471	10	Y
Chart decoder [9]	2018	11,174	5	N
Self [23]	2019	2500	2	N
Synthetic [10]	2019	202550	10	Y
UB-PMC [10]		4242	7	Y
DocFigure [21]	2019	33000	28	Y
Self [3]	2020	2702	10	N
Self [2]	2020	21099	13	N
Chart-OCR [26]	2021	386966	3	N
UB-PMC [12]	2021	22924	15	Y
Self [4]	2021	3002	10	N
Self [35]	2021	110182	24	N
UB-PMC [13]	2022	33186	15	Y

ChartSense [18]: The ChartSense dataset was put together using the ReVision dataset, and the authors manually added some additional charts. The corpus has 5659 chart images that cover ten chart categories.

ChartVega [6]: This dataset has ten chart types and was created due to a need for a benchmark dataset for chart image classification [6]. The dataset contains both synthetic and real chart images. The set contains 14,471 chart images, of which 12059 are for training and 2412 are for testing. In addition, a validation

set of 2683 real chart images is provided. No separate annotations are provided, as chart images are separated according to their types.

DocFigure [21]: This corpus consists of 28 categories of annotated figure images. There are 33,000 images that include non-chart categories like natural images, tables, 3D objects, and medical images. The train set consists of 19,797 images, and the test set contains 13173 images. The labels are provided in a text document.

ChartOCR [26]: The dataset contains 386,966 chart images created by the authors by crawling public excel sheets online. The dataset contains only three classes of chart images. The dataset is divided into the train, validation, and test sets. The training corpus contains 363,078 images, the validation set contains 11,932 images, and the test set contains 11,965 images. The annotations for the chart images are provided in JSON format.

UB-PMC CHART-Infographics: This dataset was introduced in the first edition of Competition on Harvesting Raw Tables from Infographics (ICPR 2019 CHART Infographics) [10]. This dataset has synthetic images created using matplotlib. For the testing, a large set of synthetic data and a small set of real chart images harvested from PubMedCentral¹ were used. The training set has 198,010 images, whereas the synthetic test set has 4540 images, and the real test set has 4242 images. The dataset has ten different chart categories.

The second edition of the competition [12] provided a dataset containing 22923 real chart images of 15 different chart categories in both training and testing sets. The training set has 15636 images, while the test set has 7287 images. The annotations for the chart image samples are provided in both JSON and XML formats. The dataset presented as a part of the third and most recent competition comprises 36183 images of 15 different chart categories. The training set contains 22,923 images, while the test set contains 13,260 images. Similar to the previous edition, the annotations are provided in JSON and XML formats. To the best of our knowledge, this is the largest publicly available dataset for chart image classification.

The existing classification data sets for charts are summarized in Table 3, and the composition of the publicly available datasets is reported in Table 4.

4 Deep Learning Models for Comparative Analysis

In this section, we briefly discuss prominent deep-learning models that have been used to study the performance of chart classification. We have selected two categories of deep learning models - CNN-based and Transformer-based for the comparative study. For CNN-based models, we have considered the proven state-of-the-art models for image classification on the large-scale benchmark dataset ImageNet [28] over the years. For vision transformer models, we have chosen the models that have been proven to outperform CNN-based models in computer vision tasks.

¹ <https://www.ncbi.nlm.nih.gov/pmc/>.

Table 4. Composition of publicly available datasets

Chart Type	UB-PMC [13]	DocFigure [21]	Chart-Sense [18]	Chart-OCR [26]	Chart-Vega [6]
Arc	–	–	–	–	1440
Area	308	318	509	–	1440
Block	–	1024	–	–	–
Bubble	–	339	–	–	–
Flowchart	–	1074	–	–	–
Heatmap	377	1073	–	–	–
Horizontal Bar	1421	–	–	–	–
Horizontal Interval	586	–	–	–	–
Line	13956	9022	619	122890	1440
Manhattan	256	–	–	–	–
Map	906	1078	567	–	–
Parallel Coordinate	–	–	–	–	1339
Pareto	–	311	391	–	–
Pie	433	440	568	76922	1440
Polar	–	338	–	–	–
Radar	–	309	465	–	–
Re-orderable Matrix	–	–	–	–	1440
Scatter	2597	1138	696	–	1640
Scatter-Line	3446	–	–	–	–
Sunburst	–	–	–	–	1440
Surface	283	395	–	–	–
Table	–	1899	594	–	–
Treemap	–	–	–	–	1440
Venn	206	889	693	–	–
Vertical Bar	9199	1196	557	187154	1512
Vertical Box	1538	605	–	–	–
Vertical Interval	671	–	–	–	–
Total	36183	33071	5659	386966	14471

4.1 ResNet [15]

The Deep Residual Network was introduced in 2015 and was significantly deeper than the previous deep learning networks. The motivation behind the model was to address the degradation problem: Degrading training accuracy with increasing depth of the model. The authors added shortcut connections, also known as skip connections, that perform the proposed identity mapping and are significantly easier to optimize than unreferenced mappings. Despite being deeper than the previous models, ResNet still needed to be simplified. It achieved the top-5 error of 3.57% and claimed the top position in the 2015 ILSVRC classification competition [28]. We use a 152-layer version of this Deep Residual Network called ResNet-152 for our classification problem.

4.2 Xception [8]

Xception is a re-interpretation of the inception module. The said inception module is replaced with depth-wise separable convolutions. The number of parameters in both Inception V3 and Xception is the same, so the slight performance improvement is due to the more efficient use of parameters. Xception shows

a better performance improvement than Inception V3 on the JFT dataset on the ImageNet dataset. It achieves the top five accuracy of 94.5%. Xception also shows promising results in the chart classification literature, as reported by [2] and [35].

4.3 DenseNet [17]

The Dense Convolutional Network, introduced in 2018, connects each layer in the network architecture to all other layers. This allows for the exchange of feature maps at every level and considers the same input as input gathered from all the previous layers rather than just one preceding layer. The difference between DenseNet and Resnet lies in the way that they combine features. ResNet combines features through summation, whereas DenseNet combines them through concatenation. DenseNet is easier to train due to the improved flow of gradients and other information through the network. The vanilla DenseNet has fewer parameters than the vanilla ResNet network. We used DenseNet-121 for our classification task as it was one of the best models for the chart image dataset as reported in [35].

4.4 ConvNeXt [25]

ConvNeXt model was introduced as a response to hierarchical transformers outperforming convnets in image classification tasks. Starting with a standard ResNet architecture, the model is carefully modified to adapt the specific characteristics of a typical hierarchical transformer. This resulted in a CNN-based model that matches the transformers in robustness and scalability across all benchmarks. ConvNeXt achieves a top-1 accuracy of 87.8% on ImageNet.

4.5 DeiT Transformer [36]

The authors proposed the Data Efficient Image Transformer(DeiT) with 86M parameters to make the existing vision transformer more adoptable. This convolution-free approach achieves competitive results against the existing state-of-the-art models on ImageNet. The proposed vision transformer achieved a top-1 accuracy of 85.2% on the ImageNet classification task. We use the base Base DeiT transformer for the chart classification task.

4.6 Swin Transformer [24]

A hierarchical transformer that employs shifting windows to obtain representations for vision tasks. The authors note that the hierarchical architecture provides linear computational complexity and scalability concerning image size. The limitation of self-attention calculation concerning noncoincident local windows due to the shifting windows allows for better cross-window connectivity. The qualities above contribute to the Swin transformer’s excellent performance

across computer vision tasks. It achieves 87.3% top-1 accuracy on the ImageNet-1k dataset. We perform experiments with all the 13 available Swin Transformer models and report their performance in Table 5. Furthermore we refer to the best performing Swin Transformer model as Swin-Chart in Table 6.

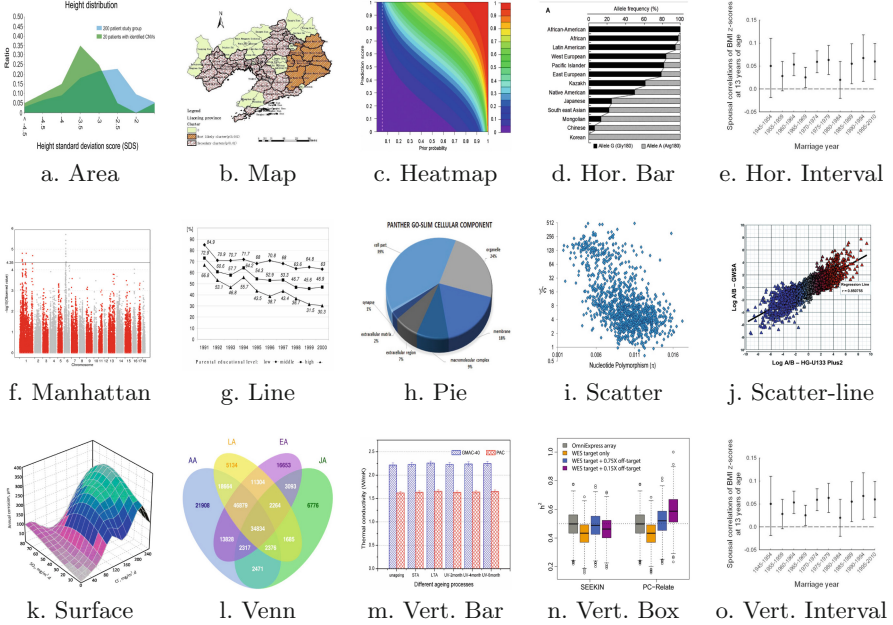


Fig. 1. Sample of chart images used in this study from UB-PMC [13] dataset

5 Experimental Protocol

5.1 Dataset

We use the ICPR2022 CHARTINFO UB PMC [13] dataset to perform our comparative study of deep learning models. The dataset is divided into training and testing sets. The number of chart images in the training and test set is 22,923 and 11,388, respectively. The ground truth values are annotated in JSON and XML formats. We further divide the provided training set into training and validation sets with an 80/20 ratio. The dataset contains charts of 15 categories: area, map, heatmap, horizontal bar, Manhattan, horizontal interval, line, pie, scatter, scatter-line, surface, Venn, vertical bar, vertical box, and vertical interval. Samples of each chart type present in the dataset are shown in Fig. 1.

5.2 Training and Testing Setup

We choose ResNet152, DenseNet121, Xception, and ConvNeXt CNN-based models and DeIT and Swin Transformers-based models for chart image classification. The CNN-based models were selected based on their performance in the existing literature on the ImageNet image classification task. The transformer-based models are chosen because they beat the CNN-based models. We use the pre-trained ImageNet weights of these models and fine-tune them for our chart classification task. The models are trained on a computer with an RTX 3090 video card with 24 GB memory. Pytorch [27] was used as the engine for our experiments. We use a batch size of 64 for CNN-based models and a batch size of 16 for transformer-based models. A learning rate of 10^{-4} is used to train each model for 100 epochs. Label Smoothing Cross Entropy Loss is used as a loss function. The evaluation measures the average over all classes and reports precision, recall, and F1-score.

Table 5. Comparative Performance of all the 13 Pre-trained Swin Transformer Models on ICPR2022 CHARTINFO UB PMC dataset

Model	Precision	Recall	F1-measure
SwinT	0.929	0.924	0.922
SwinT_s3	0.931	0.923	0.922
SwinS	0.931	0.926	0.925
SwinS_s3	0.928	0.922	0.919
SwinB_224	0.933	0.926	0.925
SwinB_384	0.936	0.932	0.931
SwinB_224_in22k_ft1k	0.934	0.930	0.929
SwinB_384_in22k_ft1k	0.933	0.929	0.927
SwinB_s3	0.927	0.923	0.921
SwinL_224	0.937	0.933	0.932
SwinL_384	0.937	0.931	0.929
SwinL_224_in22k_ft1k	0.937	0.933	0.932
SwinL_384_in22k_ft1k	0.934	0.930	0.929

5.3 Comparative Results

The models were trained following the steps mentioned in the previous section and were tested on the UB-PMC test data set. We calculate all deep learning models' average precision, recall, and F1 score. Among CNN-based models, ResNet-152 and ConvNeXt provide the best results across all evaluation metrics. The ResNet-152 result is consistent with the results in [13] for CNN-based models. For Swin transformer we perform experiments on 13 models consisting Swin Tiny(SwinT), Swin Small(SwinS), Swin Base(SwinB) and Swin Larger(SwinL)

Table 6. Comparative performances of the CNN-based and Transformer-based models on ICPR2022 CHARTINFO UB PMC dataset

Model	Precision	Recall	F1-score
Resnet-152 [15]	0.905	0.899	0.897
Xception [8]	0.882	0.870	0.866
DenseNet-121 [17]	0.887	0.879	0.875
ConvNeXt [25]	0.906	0.898	0.896
DeIT [36]	0.888	0.879	0.874
Swin-Chart	0.937	0.933	0.932

Table 7. Comparison of Swin-Chart from Table 6 with models stated in [13] on the ICPR2022 CHARTINFO UB PMC dataset

Team	Precision	Recall	F1-score
Our (Swin-Chart)	0.937	0.933	0.932
IIT_CVIT	0.926	0.901	0.910
UB-ChartAnalysis	0.900	0.881	0.886
six seven four	0.865	0.808	0.827
CLST-IITG	0.704	0.657	0.654

and their variants. SwinL with input image dimension 224 performs best with an F1-score of 0.932. So, **SwinL** model is further referred as **Swin-Chart**. The scores of all the Swin Transformer models are summarized in Table 5. The best performing CNN based models fail to compete with Swin-Chart for the chart classification task as it outperforms the other five models with an average F1-score of 0.932. The scores for the deep learning models are summarized in Table 6.

Furthermore, we compare our best-performing model(Swin-Chart) with the models reported in [13]. This comparison is summarized in Table 7. We note that Swin-Chart surpasses the winner of the ICPR 2022 CHART-Infographics competition with an average F1-score of 0.931.

6 Future Directions

Although there has been a significant increase in published articles on chart classification, several problems still need to be addressed.

6.1 Lack of Standard Benchmark Data Sets

The chart image classification problem has been extensively addressed in previous work. Efforts have been made to increase the size of chart image datasets that also cover a wide variety of charts [10, 35]. With the growing literature in various domains, authors are finding creative ways to use different charts. This

adds to the variety of chart types. Integrating such diverse chart types while creating chart datasets remains an open challenge. In addition, the popularity of charts such as bar, line, and scatter over others such as Venn, surface, and area adds to the problem of disparity between the number of samples in particular chart types.

6.2 Lack of Robust Models

Recent work makes some problematic assumptions in addressing this problem [11]. A lack of a diverse benchmark dataset adds to this problem, as there needs to be more consistency in model performance across publicly available datasets. The inherent intra-class dissimilarity and inter-class similarity of several chart types affect the model’s performance.

6.3 Inclusion of Noise

Most of the work in the existing literature ignores the effect of noise. Different types of noise, such as background grids, low image quality, composite charts, and multiple components along with figures, lead to poor performance for models that perform exceptionally well on noiseless data [34]. In addition to the noiseless chart image dataset, if a small set of chart images could be provided that incorporates the noisy images, it would help fine-tune the models to work through the inclusion of noise and be invariant to the same.

7 Conclusion

We have provided a brief survey of existing chart classification techniques and datasets. We used a Transformer model to obtain state-of-the-art results. Although there has been a significant development both in terms of variety in models and in the size of datasets, we observe that the chart classification problem still needs to be solved, especially for noisy and low-quality charts. Our comparative study showed that Swin-Chart outperforms the other vision transformer and CNN-based models on the latest UB-PMC dataset. In the future, we plan to generalize the results of the Swin-Chart over other publicly available datasets and try to bridge the gap to a robust deep-learning model for chart image classification.

References

1. Amara, J., et al.: Convolutional neural network based chart image classification. In: International Conference in Central Europe on Computer Graphics, Visualization, and Computer Vision (2017)
2. Araújo, T., et al.: A real-world approach on the problem of chart recognition using classification, detection, and perspective correction. *Sensors* **20**(16), 4370 (2020)

3. Bajić, F., et al.: Data visualization classification using simple convolutional neural network model. *Int. J. Electr. Comput. Eng. Syst. (IJECES)* **11**(1), 43–51 (2020)
4. Bajić, F., Job, J.: Chart classification using siamese CNN. *J. Imaging.* **7**, 220 (2021)
5. Balaji, A., et al.: Chart-text: a fully automated chart image descriptor. *ArXiv* (2018)
6. Chagas, P., et al.: Evaluation of convolutional neural network architectures for chart image classification. In: *International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8 (2018)
7. Cheng, B., et al.: Graphical chart classification using data fusion for integrating text and image features. In: *Proceedings of the International Conference on Document Analysis and Recognition, ICDAR* (2013)
8. Chollet, F.: Xception: deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017)
9. Dai, W., et al.: Chart decoder: generating textual and numeric information from chart images automatically. *J. Vis. Lang. Comput.* **48**, 101–109 (2018)
10. Davila, K., et al.: ICDAR competition on harvesting raw tables from infographics (CHART-infographics). In: *International Conference on Document Analysis and Recognition (ICDAR)*, pp. 1594–1599. IEEE, Sydney (2019)
11. Davila, K., et al.: Chart mining: a survey of methods for automated chart analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**(11), 3799–3819 (2021)
12. Davila, K., et al.: ICPR 2020 - competition on harvesting raw tables from infographics. In: *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event*, pp. 361–380 (2021)
13. Davila, K., et al.: ICPR: challenge on harvesting raw tables from infographics (CHART-infographics). In: *26th International Conference on Pattern Recognition (ICPR)*, pp. 4995–5001 (2022)
14. Gao, J., et al.: View: visual information extraction widget for improving chart images accessibility. In: *19th IEEE International Conference on Image Processing*, pp. 2865–2868 (2012)
15. He, K., et al.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016)
16. Howard, A.G., et al.: MobileNets: efficient convolutional neural networks for mobile vision applications (2017). <http://arxiv.org/abs/1704.04861>
17. Huang, G., et al.: Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017)
18. Jung, D., et al.: ChartSense: interactive data extraction from chart images. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017)
19. Karthikeyani, V., Nagarajan, S.: Machine learning classification algorithms to recognize chart types in portable document format (PDF) files. *IJCA* **39**(2), 1–5 (2012)
20. Krizhevsky, A., et al.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems* (2012)
21. kv, J., et al.: DocFigure: a dataset for scientific document figure classification. In: *International Conference on Document Analysis and Recognition Workshops (ICDARW)* (2019)
22. Liu, X., et al.: Chart classification by combining deep convolutional networks and deep belief networks. In: *13th International Conference on Document Analysis and Recognition (ICDAR)*, pp. 801–805 (2015)

23. Liu, X., et al.: Data extraction from charts via single deep neural network. arXiv preprint [arXiv:1906.11906](https://arxiv.org/abs/1906.11906) (2019)
24. Liu, Z., et al.: Swin transformer: hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)
25. Liu, Z., et al.: A ConvNet for the 2020s. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)
26. Luo, J., et al.: ChartOCR: data extraction from charts images via a deep hybrid framework. In: IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 1916–1924. IEEE, Waikoloa (2021)
27. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, pp. 8026–8037. Curran Associates Inc., Red Hook (2019)
28. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015)
29. Savva, M., et al.: ReVision: automated classification, analysis and redesign of chart images. In: Proceedings of the 24th annual ACM symposium on User interface software and technology, pp. 393–402. Association for Computing Machinery, New York (2011)
30. Siegel, N., Horvitz, Z., Levin, R., Divvala, S., Farhadi, A.: FigureSeer: parsing result-figures in research papers. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9911, pp. 664–680. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46478-7_41
31. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2015). <http://arxiv.org/abs/1409.1556>
32. Szegedy, C., et al.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)
33. Szegedy, C., et al.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016)
34. Thiyam, J., et al.: Challenges in chart image classification: a comparative study of different deep learning methods. In: Proceedings of the 21st ACM Symposium on Document Engineering, pp. 1–4. Association for Computing Machinery, New York (2021)
35. Thiyam, J., et al.: Chart classification: an empirical comparative study of different learning models. Presented at the December 19 (2021)
36. Touvron, H., et al.: Training data-efficient image transformers & distillation through attention. In: Proceedings of the 38th International Conference on Machine Learning, pp. 10347–10357. PMLR (2021)