

CICA: Content-Injected Contrastive Alignment for Zero-Shot Document Image Classification

Sankalp Sinha^{1,2}[0009-0009-6820-3633], Muhammad Saif Ullah Khan^{1,2}[0000-0002-7375-807X], Talha Uddin Sheikh^{1,2}[0009-0004-9156-5679], Didier Stricker^{1,2}, and Muhammad Zeshan Afzal^{1,2}[0000-0002-0536-6867]

¹ Department of Computer Science, RPTU, 67663 Kaiserslautern, Germany

² German Research Center for Artificial Intelligence (DFKI), 67663 Kaiserslautern, Germany

{sankalp.sinha, muhammad.saif_ullah.khan, talha.uddin.sheikh, didier.stricker, muhammad.zeshan.afzal}@dfki.de

Abstract. Zero-shot learning has been extensively investigated in the broader field of visual recognition, attracting significant interest recently. However, the current work on zero-shot learning in document image classification remains scarce. The existing studies either focus exclusively on zero-shot inference, or their evaluation does not align with the established criteria of zero-shot evaluation in the visual recognition domain. We provide a comprehensive document image classification analysis in Zero-Shot Learning (ZSL) and Generalized Zero-Shot Learning (GZSL) settings to address this gap. Our methodology and evaluation align with the established practices of this domain. Additionally, we propose zero-shot splits for the RVL-CDIP dataset. Furthermore, we introduce CICA³, a framework that enhances the zero-shot learning capabilities of CLIP. CICA consists of a novel 'content module' designed to leverage any generic document-related textual information. The discriminative features extracted by this module are aligned with CLIP's text and image features using a novel 'coupled-contrastive' loss. Our module improves CLIP's ZSL top-1 accuracy by 6.7% and GZSL harmonic mean by 24% on the RVL-CDIP dataset. Our module is lightweight and adds only 3.3% more parameters to CLIP. Our work sets the direction for future research in zero-shot document classification.

Keywords: Document Image Classification · Zero-Shot Learning · Generalised Zero-Shot Learning · Multimodal Learning

1 Introduction

Document image classification is the task of categorizing document images into predefined classes based on their visual and/or textual content [28]. This process enhances efficiency by automating the organization of vast quantities of digitized documents [8,28]. It significantly contributes to downstream processes

³ pronounced "ki·ka".

like information extraction by making it easier to extract valuable insights from documents [8]. Moreover, document image classification is vital for ensuring compliance and effective record-keeping across several industries, including law, healthcare, and finance [8,28]. In today’s digital era, where the volume of digitized documents is rapidly increasing, automated classification is an essential tool to manage this large data effectively [8,28].

Traditional approaches have treated document image classification as a supervised learning problem [1,2,21,16,39]. In supervised learning, the model learns to map input data to the corresponding labels, focusing on features directly relevant to accomplish this task [28]. This requires extensive labeled data for training. Recent advancements in visual recognition tasks have moved towards self-supervised learning [9,14] to mitigate this reliance on extensive labeling. Self-supervised pretraining alleviates the burden of feature extraction by utilizing auxiliary pretext tasks, which facilitate the network in comprehending the general structure of images [15]. After this pretraining phase, the model undergoes fine-tuning in a supervised manner with the available labeled dataset. A novel approach in this domain involves integrating these self-supervised pretext tasks to enhance the efficacy of document image classification systems [10].

Methods utilizing self-supervised and transfer learning significantly reduce dependence on labeled data. Nonetheless, these methods encounter challenges with domain shifts and exhibit limited zero-shot learning capabilities. This limitation is particularly evident when dynamically modifying the number of classes during runtime [23,37].

In zero-shot learning, models trained on a set of seen classes are asked to predict classes not in training data. CLIP (Contrastive Language–Image Pre-training) [34] represents a significant leap in zero-shot transfer capabilities in the visual recognition domain. It demonstrated impressive performance on multiple benchmarks by aligning images with textual descriptions. CLIP utilizes a large image-text pair dataset, enabling the model to extensively learn various visual concepts and their associations with natural language [34].

The field of zero-shot learning for document images is still emerging. However, the zero-shot proficiency of CLIP has sparked some notable research, such as the study by Khalifa et al. [23], demonstrating the enhancement of zero-shot classification through contrastive training in unstructured documents. Additionally, the analysis by Siddiqui et al. [37] offers a critical examination of CLIP’s zero-shot transfer capabilities using the RVL-CDIP dataset [17].

The current work on zero-shot learning in document images remains scarce [37,23], with existing studies lacking a unified evaluation framework. This absence of a standardized evaluation criterion is particularly notable, considering the well-established evaluation practices in the broader field of visual recognition, as meticulously outlined by Xian et al. [42]. Xian et al. [42] propose a unified evaluation framework that emphasizes average per-class top-1 accuracy as a critical metric to ensure fairness across datasets, particularly those with imbalanced class distributions, in the zero-shot learning (ZSL) scenario. Additionally, it introduces the harmonic mean of seen and unseen accuracies as a balanced eval-

uation metric for generalized zero-shot learning (GZSL) scenarios. To the best of our knowledge, zero-shot learning has not been comprehensively implemented within the document image domain. This area continues to be under-researched and represents a significant gap in current literature.

In this paper, we address this research gap of zero-shot learning in document image classification. To achieve this, we propose a novel 'content module' specifically designed to process general information from documents. In our case, this information consists of text extracted from document images using OCR. Moreover, we present a unique 'coupled contrastive' loss mechanism. This mechanism is designed to align the features from our 'content module' with both the text and image features of the CLIP model, adhering to CLIP's contrastive alignment principles. These innovations together form a cohesive, end-to-end trainable framework called CICA that significantly improves the zero-shot learning capabilities of CLIP. This framework is meticulously evaluated under the standardized criterion outlined in [42]. Our code is publicly available⁴. The primary contributions of our research are outlined as follows:

- This work represents, to the best of our knowledge, the first instance of implementing Zero-Shot Learning (ZSL) and Generalized Zero-Shot Learning (GZSL) within the domain of document image classification.
- We introduce a novel 'content module' that processes generic information from documents. In our context, this information is the OCR-extracted text from document images.
- We introduce a novel 'coupled contrastive loss' that aligns the features of our 'content module' with the text and image features of CLIP.
- We propose novel ZSL and GZSL splits for the RVL-CDIP dataset, grounded on CLIP's performance, and conduct a quantitative analysis of CLIP's zero-shot transfer potential.
- We conduct comprehensive experiments and ablation studies within the ZSL and GZSL settings to rigorously evaluate our CICA framework.

2 Related Work

2.1 Document Image Classification

Supervised deep learning approaches have significantly outperformed rule-based methods and hand-crafted features in document image classification [21,1,16,39,2,25,20], primarily due to their ability to learn complex patterns and representations directly from data. Unlike rule-based systems [29,22,3,31,27] that rely on predefined rules and manual feature engineering, deep learning models automatically extract and learn the most relevant features for a given task. However, the effectiveness of deep learning models heavily depends on the availability of large training data. This requirement can be challenging, as obtaining and labeling a vast and varied dataset of document images is time-consuming and resource-intensive [28].

⁴ <https://github.com/mindgarage/cica>

In the field of document image classification, another prevalent approach involves leveraging transfer learning [11,1] specifically fine-tuning to attain competitive results. This method involves pre-training a model on a substantial image-classification dataset, such as ImageNet [26], thereby transforming the initial layers of the convolutional network, or potentially even transformer-based models like Vision Transformers (ViT) [13] into universal feature extractors. Subsequently, the model undergoes fine-tuning, particularly after reinitializing the output layer to suit the specific dataset.

Recent trends in visual recognition have shifted towards semi-supervised and self-supervised learning to reduce reliance on extensive labeled datasets. In SimCLR [9], Chen et al. introduce a contrastive learning-based self-supervised framework to distinguish similar and dissimilar image pairs, enhancing learning without labeled data. Grill et al. present BYOL [14], which avoids negative pairs in contrastive learning and uses a dual neural network for learning from image augmentations. Following these advancements, Cosma et al. [10] introduce multimodal self-supervised learning for document images, using techniques such as 'Jigsaw Whole' for effective context-sensitive pretraining and text topic prediction from images in data-scarce scenarios. Dauphinee et al. [12] further the application of multimodal models by employing separate classifiers for visual and textual data, achieving notable accuracy on the RVL-CDIP dataset. These sequential studies highlight the evolving sophistication in multimodal integration for document image classification.

2.2 Zero-Shot Image Classification

Zero-Shot Learning (ZSL) is a machine learning paradigm designed to classify objects or entities not present in the training data. In ZSL, a model trained on a set of (seen) classes is evaluated on a different set of previously unseen classes [42]. Generalized zero-shot learning extends ZSL by including both seen and unseen classes in the test phase, making it a more realistic and challenging scenario. GZSL aims to balance the classification performance across both seen and unseen classes, addressing the inherent bias towards seen classes due to the training process [42].

In the literature, this is accomplished by utilizing supplementary information associated with a unique group of previously unknown classes. Various approaches seek to create a link between image attributes and a class representation that incorporates supplementary information [6,30,35,41]. However, these methods frequently face challenges with biases towards classes that have not been previously observed. To address this issue, a different set of techniques aims to understand the distribution of image features through generative modeling [4,40,36]. In this scenario, some methods directly utilize semantic information to produce features for new and previously unknown classes [36,44]. Certain strategies concentrate on developing a class-conditional generator capable of generating features for classes that have not been seen before [40,4]. After being trained with data from known classes, these models employ supplementary information to create features for classes they have not encountered before, thereby

mitigating issues of bias [33]. Additionally, other studies aim to improve the integration of visual and semantic data [5,19] and to enhance the performance of feature extraction from images through advanced training methods [43,18].

CLIP [34], noteworthy for its zero-shot transfer capabilities, revolutionized computer vision and NLP integration by learning from natural language without relying on labeled data. It used a vast collection of 400 million internet-sourced image-text pairs for training, focusing on matching images to text descriptions. Integrating additional modalities into CLIP to enhance its native performance in zero-shot settings is an emerging research field. Our work draws inspiration from Khan et al. [24], who demonstrate that heatmaps can be used as an additional modality within the CLIP framework to enhance the performance of various human-centric classification tasks.

2.3 Zero-Shot Document Image Classification

Zero-shot learning, within the context of document image classification, is an emerging subfield within zero-shot learning that has not been extensively studied. Khalifa et al. [23] demonstrate how contrastive training methods can enhance zero-shot classification performance for unstructured documents. Similarly, the research by Siddiqui et al. [37] focuses on investigating the potential of zero-shot learning for document image classification. Their study provides a critical analysis of the zero-shot transfer abilities of the CLIP model. It leverages CLIP’s pre-trained weights to evaluate the average top-1 accuracy across the RVL-CDIP dataset, highlighting CLIP’s zero-shot potential.

3 Methodology

Section 3 offers a comprehensive overview of our proposed CICA framework. In Subsection 3.1, we explore the problem formulation for zero-shot learning. Subsection 3.2 introduces our innovative ‘content module’ and its integration within the CLIP model. Subsection 3.3 details the formulation of our proposed novel ‘coupled contrastive’ loss. Lastly, Subsection 3.4 outlines the procedures for model inference. Together, these subsections systematically develop the narrative of the end-to-end trainable CICA framework, covering everything from its theoretical basis to its training and inference mechanisms.

3.1 Problem Formulation

Zero-shot learning aims to enhance a model’s ability to generalize to classes not present in its initial training. This is achieved by leveraging auxiliary information related to a distinct set of unknown classes. In this context, we define the set of classes in the training dataset as Y_s and the set of classes only encountered during testing as Y_u . Hence, Y_s is the set of seen classes, and Y_u is the set of unseen classes. Consider our training dataset $T = \{(x, y, c) | x \in X_s, y \in Y_s, C \in C_s\}$. In this dataset, x is an RGB image from the training set X_s . The term y refers to

the corresponding label from the seen classes Y_s . Furthermore, c represents the associated information (content) from the document, where C_s is the content set for the seen classes. During testing, an additional content set C_u is also made available, corresponding to the novel, unseen classes. This is defined such that $C = C_s \cup C_u$ and $C_s \cap C_u = \emptyset$. In the ZSL scenario, the model has to classify an instance into one of the unseen classes Y_u . In the GZSL scenario, the model’s objective is extended to accurately classify instances into both the seen and unseen classes, combined as $Y = Y_s \cup Y_u$.

3.2 Model Architecture

The proposed model extends the CLIP architecture by integrating a plug-and-play ‘content module’ as seen in *Fig. 1*. This module is trained using our proposed ‘coupled contrastive’ loss, which harmonizes the features of the content encoder with those of the image and text encoders.

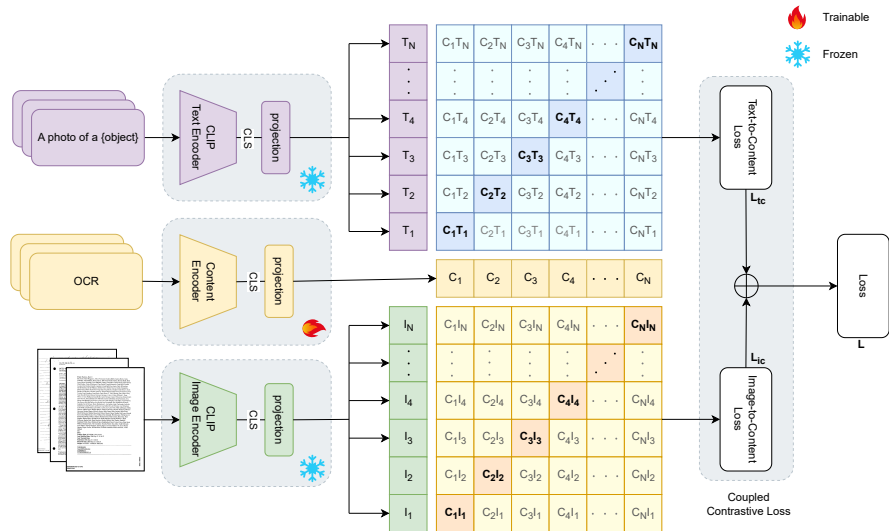


Fig. 1. Illustration of CICA’s architecture and training paradigm, highlighting how image and text features from CLIP are integrated with the content module’s features through a coupled contrastive loss. This aligns positive image-content and text-content pairs (highlighted in the matrix) and distances negative pairs (lightly shaded in the matrix). N denotes the batch size.

Our model leverages CLIP’s image and text encoders to process both visual and textual data. Our proposed content module includes a transformer encoder block, complete with a dedicated token and position embedding layer.

The ‘coupled contrastive’ loss introduces an element of self-supervised learning. It promotes a unified and coherent feature space among all encoders in the model, empowering the content encoder to explore and delineate inherent relationships within the data autonomously. This strategy aligns seamlessly with the contrastive pre-training framework employed by CLIP.

3.3 Loss Function

Given a batch of N images $\mathbb{I} = \{x_i\}$, texts $\mathbb{T} = \{t_i\}$, and contents $\mathbb{C} = \{c_i\}$, let f_i , f_t , and f_c represent the image, text, and content encoders, respectively. Each image $x_i \in \mathbb{R}^{H \times W \times Ch}$, where H and W denote the image’s height and width, and Ch represents the RGB color channels. The images are converted into 2D flattened patches $p_i \in \mathbb{R}^{N \times (P^2 \cdot Ch)}$, where P is the patch size. This process results in $M = \frac{HW}{P^2}$ patches per image. Each text t_i is tokenized using CLIP’s tokenizer, with a context length of 77. This tokenization yields a sequence of tokens denoted by t'_i . Similarly, each content c_i is tokenized with a context length of 512, producing a sequence denoted by c'_i . For each sequence namely: tokenized text t'_i , tokenized content c'_i , and flattened image patches p_i , a CLS token is prepended. The resulting sequences are denoted as $[CLS]_i^I$, $[CLS]_i^T$, and $[CLS]_i^C$ respectively. These sequences are then passed through their respective encoders: f_i , f_t , and f_c . After processing, the CLS tokens are extracted and represented as CLS_i^I , CLS_i^T , and CLS_i^C . Finally, these CLS tokens are projected into a joint embedding space of 512 dimensions using projection layers. These layers are denoted as Proj^I , Proj^T , and Proj^C as seen in Equation 1. The resulting features are normalized, as detailed in Equations 2.

$$\hat{I}_i = \text{Proj}^I(CLS_i^I) \quad \hat{T}_i = \text{Proj}^T(CLS_i^T) \quad \hat{C}_i = \text{Proj}^C(CLS_i^C) \quad (1)$$

$$I_i = \frac{\hat{I}_i}{\|\hat{I}_i\|} \quad T_i = \frac{\hat{T}_i}{\|\hat{T}_i\|} \quad C_i = \frac{\hat{C}_i}{\|\hat{C}_i\|} \quad (2)$$

$$L_{i,ic} = -\log \frac{\exp(\phi(I_i, C_i)/\tau_{IC})}{\sum_{j=1}^N \mathbb{1}_{[j \neq i]} \exp(\phi(I_i, C_j)/\tau_{IC})} \quad (3)$$

$$L_{i,tc} = -\log \frac{\exp(\phi(T_i, C_i)/\tau_{TC})}{\sum_{j=1}^N \mathbb{1}_{[j \neq i]} \exp(\phi(T_i, C_j)/\tau_{TC})} \quad (4)$$

Eq. 3 and Eq. 4 denote the image-to-content and text-to-content loss components, where τ_{IC} and τ_{TC} are the learned temperature parameters for each component. Furthermore, ϕ represents cosine similarity, and $\mathbb{1}$ is an indicator function. Eq. 5 represents the total loss as the average contrastive loss for each image-text-content triplet in the batch, combining the losses from both image-to-content and text-to-content pairs.

$$L = \sum_{i=1}^N \frac{1}{2} (L_{i,ic} + L_{i,tc}) \quad (5)$$

3.4 Inference

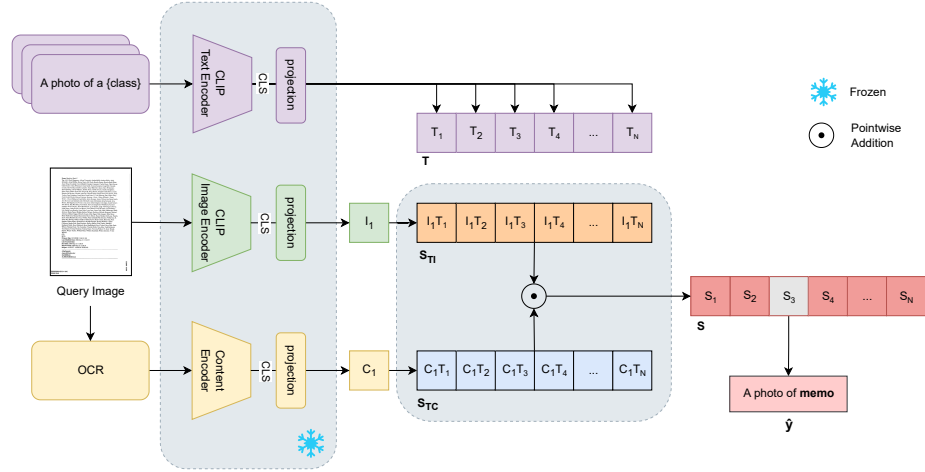


Fig. 2. Illustration showing the inference logic for CICA. At test time, the learned content encoder, along with the CLIP text encoder, synthesizes a zero-shot linear classifier by embedding a prompt with the names of the test set’s classes along with the content for the test sample. Here, N is the number of classes.

Given a test image $x_i \in \mathbb{I}$ and the content for the test image $c_i \in \mathbb{C}$. Let $\{t_i\}$ be the set of texts for each class. Using Eq. 2, we get the normalized feature I_1 and C_1 for the image and the content, along with the set of normalized features for the texts denoted as $T = \{T_i\}$ where $i \in [1, N]$ and N is the number of classes. Now, Eq. 6 and Eq. 7 represent the text-to-image and text-to-content logits. The final logits and prediction are represented in Eq. 8 and Eq. 9, respectively.

$$S_{TI} = \tau_{IC} \times I_1^T \cdot T \quad (6)$$

$$S_{TC} = \tau_{TC} \times C_1^T \cdot T \quad (7)$$

$$S = \frac{S_{TI} + S_{TC}}{2} \quad (8)$$

$$\hat{y} = \arg \max S \quad (9)$$

4 Experimental Setup

This section details the experimental setup for evaluating the proposed CICA framework. In Subsection 4.1, we describe the dataset utilized. In Subsection 4.2, we

outline the evaluation criterion. In Subsection 4.3, we present the proposed zero-shot data splits. Finally, in Subsection 4.5, we provide implementation details of the model.

4.1 Datasets

RVL-CDIP The RVL-CDIP dataset [17] comprises 400,000 grayscale images from a broader collection of 1 million images. It is organized into 16 classes: letters, forms, emails, handwritten notes, advertisements, scientific reports, publications, specifications, file folders, news articles, budgets, invoices, presentations, questionnaires, resumes, and memos. Its wide variety and substantial volume make it a valuable resource for training and evaluating machine learning models. It is especially instrumental for document classification and OCR tasks.

4.2 Evaluation Criteria

We follow the ZSL and GZSL protocols proposed by Xian et al. [42]. Single-label image classification accuracy is typically assessed using Top-1 accuracy. A prediction is considered accurate if the top predicted class matches the correct one. This accuracy, when averaged across all images, tends to favor densely populated classes. To achieve high performance across both densely and sparsely populated classes, we adjust our method. We start by independently calculating the accuracy of each class. Then, we compute their cumulative sum. This is divided by the total number of classes, effectively measuring the average per-class Top-1 accuracy. This method offers a more balanced evaluation across classes of varying sizes and is used in our ZSL setting [42].

$$\text{acc}_Y = \frac{1}{|Y|} \sum_{n=1}^{|Y|} \frac{\# \text{ correct predictions in } n}{\# \text{ samples in } n} \quad (10)$$

The GZSL setting evaluates both test classes (Y_u) and training classes (Y_s). This makes the approach more relevant to real-world scenarios. In our testing procedure, we utilize a subset of images from the training classes to calculate the average per-class top-1 accuracy for both the training and test classes. We then compute the harmonic mean of these accuracies, offering a comprehensive measure of the model’s performance across seen and unseen classes [42]. Additionally, we report the individual accuracies for both seen and unseen classes.

$$H = 2 \times \frac{\text{acc}_{Y_s} \times \text{acc}_{Y_u}}{\text{acc}_{Y_s} + \text{acc}_{Y_u}} \quad (11)$$

Calibrated stacking [7] is widely used in the GZSL setting. We also employ it for our evaluations. Our evaluation methodology involves utilizing a pre-trained and unaltered CLIP model, as described in Radford et al. [34]. We conduct comparative assessments of both CLIP and our proposed CICA model. These assessments occur in ZSL and GZSL settings. We adhere to the evaluation criteria outlined in section 4.2. We also utilize the data splits outlined in section 4.3.

4.3 Splits for Zero-Shot Learning

We propose two distinct dataset splits for the RVL-CDIP dataset for both the ZSL and GZSL settings. The first is the ‘sequential splits,’ which designate four classes as unseen while keeping the remaining twelve as seen. We create four such distinct splits, covering all 16 dataset classes. This approach guarantees that each class is assigned to the unseen set exactly once. *Tab. 1* illustrates these splits.

The second is the ‘incremental splits,’ produced using rank ordering. Given a rank ordering R of the top-1 class-wise accuracies from the RVL-CDIP dataset in ascending order, obtained using the pre-trained CLIP model. We define the split S_i as follows:

$$S_i^I = (u, s)$$

where s represents the set of seen classes and u represents the set of unseen classes. In this context, the index i ranges from 2 to 8, indicating the number of classes in the unseen set u . The remaining classes, totaling $16 - i$, are included in the seen set s . In accordance with our evaluation criterion in Section 4.2, we ensure that $Y_u \cap Y_s = \emptyset$ and $Y = Y_u \cup Y_s$. The rank-ordering R for the RVL-CDIP dataset is shown in *Fig. 3*. Essentially, S_2^I keeps the top two worst classes in the unseen set and the rest in the seen set. Similarly, S_3^I keeps the top three worst performing classes in the unseen set and so on for S_4^I to S_8^I .

Table 1. Sequential splits of the RVL-CDIP dataset where four classes are kept in the unseen set, with the rest in the seen set. Table shows the unseen classes for each split.

Split Unseen Classes	
A	email, form, handwritten, letter
B	advertisement, scientific publication, scientific report, specification
C	budget, file folder, invoice, news article
D	memo, presentation, questionnaire, resume

4.4 Document Information or Content

Our content module is designed to be adaptable, allowing it to process a wide range of additional information about the target classes. However, in our studies, we focus on using text extracted through Optical Character Recognition (OCR) as the content. We employ two OCR engines, Tesseract [38] and docTR [32], for this purpose, which are crucial in creating the content set C. The impact of these two OCR engines on our model’s performance is analyzed in *Tab. 4*.

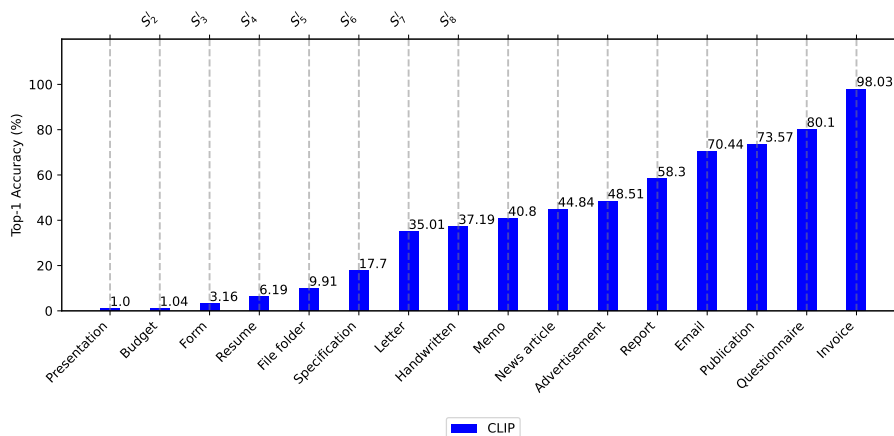


Fig. 3. Rank ordering of class-wise top-1 accuracies obtained from the frozen CLIP model, for the RVL-CDIP dataset. The top axis also shows the zone of the varying number of unseen classes in the incremental splits (S_i^I).

4.5 Implementation Details

Our CICA model was developed with PyTorch, mirroring CLIP’s data augmentation techniques. We set the learning rate at $1e-3$, reducing it to 0.1 after five epochs. The model was trained with a batch size of 32 over 10 epochs. For optimization, we used the Adam optimizer. Unless otherwise specified, all models utilize the OCR as content from the docTR OCR engine. Our model uses the ViT-L/14@336px variant of CLIP as its backbone architecture. All comparisons are against this specific CLIP version. For the finetuned version of CLIP, we use a learning rate at $1e-4$, reducing it to 0.1 after five epochs and training for 10 epochs in total. Content for the 400,000 grayscale images in the RVL-CDIP dataset was generated using two OCR engines in ≈ 24 hours with 64 CPUs.

5 Results and Discussion

5.1 Results

Tab. 2 provides a detailed comparison of the performance between the different models. The table focuses explicitly on the sequential splits of the RVL-CDIP dataset. The results show that our CICA model consistently surpasses CLIP in performance across the splits. While a substantial increase in accuracies for seen classes is observed, it can be noted that simple finetuning of CLIP does not result in performance gains in ZSL and GZSL settings. On the other hand, we observe an average increase of 6.7% in the top-1 accuracy across the splits for the ZSL setting. Further, in the GZSL setting, we see an average improvement of 24% in the harmonic mean (**H**).

Table 2. Comparative performance of CLIP, finetuned CLIP, and CICA across sequential splits (A - D) of the RVL-CDIP dataset, highlighting top-1 accuracy and harmonic mean in **bold**. CICA consistently outperforms CLIP in both ZSL and GZSL settings on all splits.

Split Method		RVL-CDIP			
		ZSL		GZSL	
		T1	u	s	H
A	CLIP (finetuned)	42.62	16.50	61.60	26.02
	CLIP	70.55	49.74	35.61	41.50
	CICA	70.64	61.84	69.36	65.38
B	CLIP (finetuned)	45.73	23.73	70.92	35.56
	CLIP	67.02	53.83	33.88	41.58
	CICA	77.87	69.84	75.40	72.52
C	CLIP (finetuned)	48.06	29.38	67.98	41.03
	CLIP	50.35	41.61	37.31	39.34
	CICA	59.52	53.94	72.17	61.74
D	CLIP (finetuned)	41.42	13.25	78.99	22.69
	CLIP	51.67	39.37	36.29	37.77
	CICA	61.13	51.13	71.17	59.51

These improvements highlight the efficacy of the CICA model. It capitalizes on the inherent capabilities of CLIP’s architecture. Additionally, it successfully integrates features tailored to the ZSL and GZSL requirements.

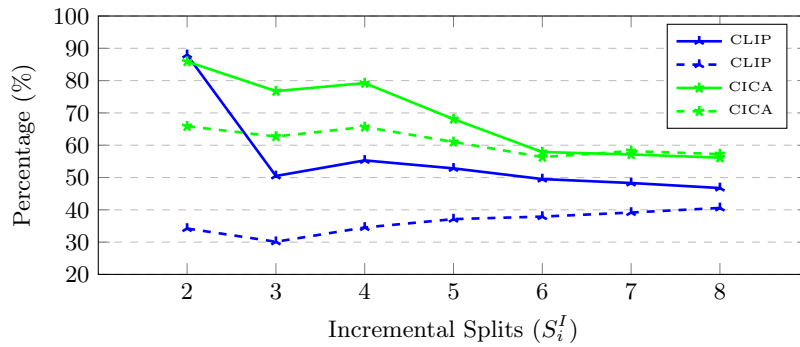


Fig. 4. Comparative analysis of CLIP and CICA on RVL-CDIP incremental splits, showing CICA outperforming CLIP. Dashed lines indicate GZSL outcomes as harmonic mean, while solid lines depict ZSL performance using Top-1 accuracy. We observe that CICA consistently outperforms CLIP.

Furthermore, *Fig. 4* illustrates the performance comparison between CLIP and our proposed CICA. The figure specifically focuses on the incremental splits of the RVL-CDIP for both the ZSL and GZSL settings. It is evident that for ZSL scenarios (solid lines), CICA outperforms CLIP in most splits except when $s = 2$, where its performance is slightly lower. This can be attributed to the 'presentation' and 'budget' classes being designated as unseen during this split. These classes are significantly different from the rest of the dataset, making it unlikely for our model to infer knowledge about them from the seen classes. In GZSL scenarios (dotted lines), both models demonstrate a decline in performance as the number of unseen classes increases, yet CICA maintains a consistent lead over CLIP. This indicates that CICA has a better capability to generalize to unseen classes. Despite the unique challenge posed by the distinct 'presentation' and 'budget' categories, we see an overall improvement in both the ZSL and GZSL results.

5.2 Ablation Study

In our ablation studies, we focus on three crucial aspects that determine our model's effectiveness. Firstly, we analyze the impact of aligning features from our content module specifically with either the image or text features of CLIP. Secondly, we evaluate how the quality of text extracted by different OCR engines influences CICA's accuracy. Thirdly, we investigate the effects of feature fusion techniques on the CICA's ability to predict class labels precisely.

Influence of feature alignment The alignment ablation studies on the CICA model aim to assess its effectiveness. We compare aligning the OCR with image features alone (C2I) or text features alone (C2T) against aligning with both modalities. The results, as detailed in *Tab. 3*, show varying degrees of performance across the four sequential splits (A to D) for the RVL-CDIP dataset. Within the ZSL setting, the CICA model's performance does not favor a single alignment strategy; both C2T and C2I alignments prove crucial across various splits. Specifically, C2T alignment is vital for splits A and B, whereas C2I alignment is key for splits C and D. For GZSL, the combined alignment of CICA consistently outperforms the singular alignment models. This affirms the importance of multimodal integration for enhanced model performance. Overall, the combined alignment approach of the CICA model demonstrates the most balanced performance, particularly in GZSL scenarios, where it achieves higher harmonic mean (**H**), showcasing its superior capability to generalize across both seen and unseen classes.

Influence of OCR engine In our OCR ablation studies for the CICA model, we examine the impact of different OCR engines on the model's efficacy. We compare Tesseract OCR with the default docTR OCR engine. *Tab. 4* shows our findings. In the ZSL setting, the performance boost with docTR is not uniform across all evaluated splits. This is evident from the top-1 accuracy rates. The

Table 3. Performance outcomes of CICA on sequential splits (A - D) of the RVL-CDIP dataset, demonstrating alignment ablations: 'C2I' for sole Content-to-Image alignment and 'C2T' for sole Content-to-Text alignment. We observe that aligning content to both text and image improves the performance of CICA. Best values are **bold**.

Split Method		RVL-CDIP			
		ZSL		GZSL	
		T1	u	s	H
A	CICA (C2I)	68.25	55.21	49.85	52.39
	CICA (C2T)	68.85	57.37	69.56	62.88
	CICA	70.64	61.84	69.36	65.38
B	CICA (C2I)	75.63	65.25	50.80	57.13
	CICA (C2T)	77.26	65.13	75.28	69.84
	CICA	77.87	69.84	75.40	72.52
C	CICA (C2I)	63.94	55.94	50.79	53.24
	CICA (C2T)	51.65	45.76	69.55	55.21
	CICA	59.52	53.94	72.17	61.74
D	CICA (C2I)	59.28	50.21	46.47	48.27
	CICA (C2T)	54.54	45.92	66.47	54.31
	CICA	61.13	51.13	71.17	59.51

situation changes in the GZSL setting. Here, docTR significantly enhances accuracies for both unseen and seen classes. It also achieves a better harmonic mean (**H**). These results highlight the crucial role of OCR quality. It improves model accuracy for both familiar and new classes. docTR’s accuracy in providing OCR greatly benefits CICA’s predictive performance.

Table 4. Performance with OCR engine variations on sequential splits (A - D) of RVL-CDIP. We observe CICA’s sensitivity to OCR quality. Best values are **bold**.

Split Method		RVL-CDIP			
		ZSL		GZSL	
		T1	u	s	H
A	CICA (Tesseract)	64.10	54.70	51.53	53.07
	CICA (docTR)	70.64	61.84	69.36	65.38
B	CICA (Tesseract)	67.88	57.13	59.61	58.35
	CICA (docTR)	77.87	69.84	75.40	72.52
C	CICA (Tesseract)	66.77	52.12	57.92	54.87
	CICA (docTR)	59.52	53.94	72.17	61.74
D	CICA (Tesseract)	65.48	52.17	51.38	51.77
	CICA (docTR)	61.13	51.13	71.17	59.51

Influence of inference strategy The inference ablation studies for CICA examine the effects of different fusion techniques on predicting class labels. *Tab. 5* outlines these comparisons. The standard method employs late fusion, first integrating content encoding C_1 and image encoding I_1 with text encoding T_1 . Subsequently, it calculates the mean of these dot products, as depicted in *Fig. 2*. Conversely, the early fusion method merges I_1 and C_1 before the dot product with T_1 . Results indicate that CICA’s late fusion performs comparably to early fusion in ZSL scenarios, with minor differences observed in splits ‘B’ and ‘C’. In GZSL scenarios, however, late fusion surpasses early fusion across all splits. This suggests that processing image and content data separately retains detailed features, improving accuracy. The harmonic mean (**H**) is consistently higher without early fusion. This demonstrates that this approach better supports generalization across both seen and unseen classes.

Table 5. CICA’s inference fusion ablation outcomes across sequential splits (A - D) of RVL-CDIP. Early fusion leads to poor GZSL performance. Best values are **bold**.

Split Method		RVL-CDIP			
		ZSL	GZSL		
		T1	u	s	H
A	CICA (Early Fusion)	70.63	65.86	59.22	62.36
	CICA	70.64	61.84	69.36	65.38
B	CICA (Early Fusion)	77.88	76.34	42.01	54.20
	CICA	77.87	69.84	75.40	72.52
C	CICA (Early Fusion)	59.55	59.01	45.25	51.22
	CICA	59.52	53.94	72.17	61.74
D	CICA (Early Fusion)	61.11	56.32	53.42	54.83
	CICA	61.13	51.13	71.17	59.51

6 Conclusion

In our study, we introduce CICA, a novel approach for zero-shot learning in document image classification. We establish dataset splits to provide a baseline for future research and incorporate a ‘content module’ for leveraging OCR data. This approach can also integrate other textual information, such as metadata. Our innovative ‘coupled contrastive’ loss mechanism aligns the features of the content module with those of CLIP, enhancing the model’s performance. Specifically, our method improves CLIP’s ZSL top-1 accuracy by 6.7% and GZSL harmonic mean by 24% on the RVL-CDIP dataset. These improvements, coupled with our development of specific ZSL and GZSL dataset splits, highlight CICA’s effectiveness and potential to advance document management systems by improving model generalization across unseen document types.

References

1. Afzal, M.Z., Capobianco, S., Malik, M.I., Marinai, S., Breuel, T.M., Dengel, A., Liwicki, M.: Deepdocclassifier: Document classification with deep convolutional neural network. In: 2015 13th international conference on document analysis and recognition (ICDAR). pp. 1111–1115. IEEE (2015)
2. Afzal, M.Z., Kölsch, A., Ahmed, S., Liwicki, M.: Cutting the error by half: Investigation of very deep cnn and advanced training strategies for document image classification. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR). vol. 1, pp. 883–888. IEEE (2017)
3. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006. Proceedings, Part I 9. pp. 404–417. Springer (2006)
4. Brendel, W., Bethge, M.: Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. arXiv preprint arXiv:1904.00760 (2019)
5. Cacheux, Y.L., Borgne, H.L., Crucianu, M.: Modeling inter and intra-class relations in the triplet loss for zero-shot learning. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10333–10342 (2019)
6. Changpinyo, S., Chao, W.L., Gong, B., Sha, F.: Synthesized classifiers for zero-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5327–5336 (2016)
7. Chao, W.L., Changpinyo, S., Gong, B., Sha, F.: An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14. pp. 52–68. Springer (2016)
8. Chen, N., Blostein, D.: A survey of document image classification: problem statement, classifier architecture and performance evaluation. *International Journal of Document Analysis and Recognition (IJDAR)* **10**, 1–16 (2007)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
10. Cosma, A., Ghidoveanu, M., Panaitescu-Liess, M., Popescu, M.: Self-supervised representation learning on document images. In: Document Analysis Systems: 14th IAPR International Workshop, DAS 2020, Wuhan, China, July 26–29, 2020, Proceedings 14. pp. 103–117. Springer (2020)
11. Das, A., Roy, S., Bhattacharya, U., Parui, S.K.: Document image classification with intra-domain transfer learning and stacked generalization of deep convolutional neural networks. In: 2018 24th international conference on pattern recognition (ICPR). pp. 3180–3185. IEEE (2018)
12. Dauphinee, T., Patel, N., Rashidi, M.: Modular multimodal architecture for document classification. arXiv preprint arXiv:1912.04376 (2019)
13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
14. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems* **33**, 21271–21284 (2020)

15. Gui, J., Chen, T., Cao, Q., Sun, Z., Luo, H., Tao, D.: A survey of self-supervised learning from multiple perspectives: Algorithms, theory, applications and future trends. arXiv preprint arXiv:2301.05712 (2023)
16. Harley, A.W., Ufkes, A., Derpanis, K.G.: Evaluation of deep convolutional nets for document image classification and retrieval. In: 2015 13th International Conference on Document Analysis and Recognition (ICDAR). pp. 991–995. IEEE (2015)
17. Harley, A.W., Ufkes, A., Derpanis, K.G.: Evaluation of deep convolutional nets for document image classification and retrieval. In: International Conference on Document Analysis and Recognition (ICDAR) (2015)
18. Ji, Z., Fu, Y., Guo, J., Pang, Y., Zhang, Z.M., et al.: Stacked semantics-guided attention model for fine-grained zero-shot learning. *Advances in neural information processing systems* **31** (2018)
19. Jiang, H., Wang, R., Shan, S., Chen, X.: Transferable contrastive network for generalized zero-shot learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9765–9774 (2019)
20. Kanchi, S., Pagani, A., Mokayed, H., Liwicki, M., Stricker, D., Afzal, M.Z.: Emm-docclassifier: Efficient multimodal document image classifier for scarce data. *Applied Sciences* **12**(3), 1457 (2022)
21. Kang, L., Kumar, J., Ye, P., Li, Y., Doermann, D.: Convolutional neural networks for document image classification. In: 2014 22nd international conference on pattern recognition. pp. 3168–3172. IEEE (2014)
22. Ke, Y., Sukthankar, R.: Pca-sift: A more distinctive representation for local image descriptors. In: Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004. vol. 2, pp. II–II. IEEE (2004)
23. Khalifa, M., Vyas, Y., Wang, S., Horwood, G., Mallya, S., Ballesteros, M.: Contrastive training improves zero-shot classification of semi-structured documents. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Findings of the Association for Computational Linguistics: ACL 2023. pp. 7499–7508. Association for Computational Linguistics, Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.findings-acl.473>, <https://aclanthology.org/2023.findings-acl.473>
24. Khan, M.S.U., Naeem, M.F., Tombari, F., Van Gool, L., Stricker, D., Afzal, M.Z.: Focusclip: Multimodal subject-level guidance for zero-shot transfer in human-centric tasks. arXiv preprint arXiv:2403.06904 (2024)
25. Kölsch, A., Afzal, M.Z., Ebbecke, M., Liwicki, M.: Real-time document image classification using deep cnn and extreme learning machines. In: 2017 14th IAPR international conference on document analysis and recognition (ICDAR). vol. 1, pp. 1318–1323. IEEE (2017)
26. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25** (2012)
27. Kumar, J., Ye, P., Doermann, D.: Learning document structure for retrieval and classification. In: Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012). pp. 1558–1561. IEEE (2012)
28. Liu, L., Wang, Z., Qiu, T., Chen, Q., Lu, Y., Suen, C.Y.: Document image classification: Progress over two decades. *Neurocomputing* **453**, 223–240 (2021)
29. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International journal of computer vision* **60**, 91–110 (2004)

30. Mancini, M., Naeem, M.F., Xian, Y., Akata, Z.: Open world compositional zero-shot learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5222–5230 (2021)
31. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *IEEE transactions on pattern analysis and machine intelligence* **27**(10), 1615–1630 (2005)
32. Mindee: doctr: Document text recognition. <https://github.com/mindee/doctr> (2021)
33. Naeem, M.F., Khan, M.G.Z.A., Xian, Y., Afzal, M.Z., Stricker, D., Van Gool, L., Tombari, F.: I2mvformer: Large language model generated multi-view document supervision for zero-shot image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15169–15179 (2023)
34. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
35. Romera-Paredes, B., Torr, P.: An embarrassingly simple approach to zero-shot learning. In: International conference on machine learning. pp. 2152–2161. PMLR (2015)
36. Schönfeld, E., Ebrahimi, S., Sinha, S., Darrell, T., Akata, Z.: Generalized zero-and few-shot learning via aligned variational autoencoders. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2019) (2019)
37. Siddiqui, S.A., Dengel, A., Ahmed, S.: Analyzing the potential of zero-shot recognition for document image classification. In: International Conference on Document Analysis and Recognition. pp. 293–304. Springer (2021)
38. Smith, R.: Tesseract OCR. <https://github.com/tesseract-ocr/tesseract> (2016), accessed: 2023-01-05
39. Tensmeyer, C., Martinez, T.: Analysis of convolutional neural networks for document image classification. In: 2017 14th IAPR international conference on document analysis and recognition (ICDAR). vol. 1, pp. 388–393. IEEE (2017)
40. Verma, V.K., Arora, G., Mishra, A., Rai, P.: Generalized zero-shot learning via synthesized examples. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4281–4289 (2018)
41. Xian, Y., Akata, Z., Sharma, G., Nguyen, Q., Hein, M., Schiele, B.: Latent embeddings for zero-shot classification. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 69–77 (2016)
42. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence* **41**(9), 2251–2265 (2018)
43. Xu, W., Xian, Y., Wang, J., Schiele, B., Akata, Z.: Attribute prototype network for zero-shot learning. *Advances in Neural Information Processing Systems* **33**, 21969–21980 (2020)
44. Zhu, Y., Elhoseiny, M., Liu, B., Peng, X., Elgammal, A.: A generative adversarial approach for zero-shot learning from noisy texts. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1004–1013 (2018)