# Classifying Books by Genre Based on Cover

Rajasree Jayaram, Harshitha Mallappa, Pavithra S, Munshira Noor B, Bhanushree K J

*Abstract*: *A book cover can convey a lot about the content of the book. Despite the adage to not evaluate something based on outward appearances, we apply machine learning to see if we can, in fact, judge a book by its cover, or more specifically by its cover art and text. The classification was done considering three different aspects - cover image only, cover text only and both image and text in a multimodal approach. Image classification was done using transfer learning with Inception-v3. For text detection from the cover image, images were first converted to greyscale and different thresholds were applied to detect maximum text. This text was then vectorized and used to train a Multinomial Naïve Bayes model. We also trained custom CNNs for image and text modalities. For multimodal classification, we examine late fusion model, where the modalities are combined at decision level, and early fusion model, where the modalities are combined at the feature level. Our results show that the late fusion model performs best in our setting. We also observe that text is more informative with respect to genre prediction and that significant efforts need to be devoted to solve this image-based classification task to a satisfactory level. This research can be used to aid product design process by revealing underlying information. It could also be used in recommender systems and to help in promotion and sales processes for automatic genre suggestion.*

*Keywords*: *Text classification, Image classification, Multimodal classification, Deep Learning, Genre Prediction.*

## I. INTRODUCTION

Since the invention of the printing press in the 15th century, books have become a widespread method of retaining and sharing information. Even with the recent trend of electronic devices, the practice of reading continues to be practised and encouraged. Every day, thousands of books are published worldwide, which means it is a daunting task for a new book to be noticed and acquire a significant readership.

Visual design gives significant impressions to transmit

facts that enhance product value. Similarly, book cover design gives an impression of the content of the book. To encourage people to buy their books, authors and publishers go to some lengths to ensure that their products are visually attractive. When a person considers reading a new book, they usually glance at the book cover and tries to guess what kind of book it is. We explore whether machine learning can be used to analyze the cover and get a reasonably accurate prediction of which genre the book belongs to.

The combination of multimodal features for the recognition and classification of visual data is an emerging method for understanding image content in recent years. The multimodal data are often correlated and complementary to each other in visual understanding, since they originate from a common source. A book cover may include images along with text such as title, author, a one-line description etc. This matter, for the most part, gives obvious hints as to the type of content within. However, such hints are sometimes understood only by the seasoned reader. This work attempts to classify books by genre based on the cover image and any text present on it. It also examines to what extent the cover of a book can determine the genre it belongs to. This classification might be valuable for organizations which provide literary services, such as libraries and online bookstores. It may also be used for recommendation system purposes. The categorization of literary works is remarkably different from conventional text classification. A title is usually a very short description and an incomplete sentence. A classifier that learns from such short text may need to be designed differently from a typical text classifier.

Machine learning is being increasingly used in different fields for a variety of purposes. It is making increasingly accurate image recognition algorithms possible. Convolutional Neural Networks are multi-layer neural networks that employ learned convolutional kernels, also called filters, as a method of feature extraction. The idea here is to use learned features rather than pre-designed features as the feature representation for image recognition since it is difficult to determine precisely what features we wish to learn about. CNNs combine multiple convolutional layers along with fully-connected layers. By increasing the depth of the network, higher-level features can be learned and discriminative parts of the images are exaggerated [4]. These deep CNNs have had successes in many fields including digit recognition in [3], [5] and large-scale image recognition in [6], [7]. The purpose of fine-grained classification is to assign labels to images having subtle differences in visual appearance such as animal species and product types. This is quite challenging as pertinent differences that are not easily apparent may not be accurately extricated by a typical classification model. Scene text, which appears in natural scenes such as road signs, product packages, and shop displays,

often possesses rich and precise meanings that are highly related to the semantics of the object or scene in the same image. The text present on a book cover provides similar information of equivalent importance to the cover image. With the recent developments in text detection and recognition described in [33], [34], it has been shown that textual cues are very beneficial to fine-grained classification.

Unification methods for multimodal learning are divided into late fusion and early fusion techniques. In late fusion information from the modalities is combined at the decision level. The general idea for this is to use separate classifiers for each modality and combine the output probabilities from these networks based on some kind of policy. In early fusion, the modalities are instead combined at the feature level. The modalities are also processed by separate models, where each model produces a vector representation for its modality. These abstract representations are then fused and put though a multimodal network to generate predictions.

In this paper, we apply machine learning to book genre prediction based on cover images. We also discuss how to efficiently integrate visual and textual cues for fine-grained image classification with CNNs. We examine the outcomes of unimodal classification – using Multinomial Naïve Bayes and a custom CNN for text classification, Inception-v3 [6] transfer learning and a custom-designed CNN for image classification - and compare the performance of different multimodal techniques. We describe network architecture and results.

## II. RELATED WORK

The existing work for automatic classification of books has been done considering a variety of attributes such as cover, book metadata as well as the entire content of the book itself. Considering only image modality, two different CNN architectures, LeNet [3] and AlexNet [20] were evaluated in [2], with AlexNet achieving higher Top 1 accuracy of 24.7%. It was also attempted to visualize what objects and text regions are being detected on the cover images using SSD and EAST detector [19]. The different fonts used on covers were analyzed in [10] and identified trends in terms of how typographic design represents the impression and the atmosphere of the content genres. Reference [8] had previously considered text modality in the form of title text, those results are compared with two neural networks – one being a shallow network, the other deeper with architecture similar to VGG16 – which they used to classify cover images. A subset of Amazon co-purchase dataset was examined in [25] to specifically examine book co-purchases and extended it to include visual information from the covers, such as dominant colours and objects detected using Darknet framework. Analysis of book covers was performed using the JRip algorithm implemented in Weka. We have also looked at work done on genre classification of comic books and movie genre based on posters since these endeavours are analogous to our purposes. Machine learning has been used in a variety of forms and combinations, to good effect. For comic books, attempts have been made to distinguish them by genre as well as artistic style. Research has also been done as to how to accurately text present in a comic book can be detected using an off-the-shelf neural network [16] – [18].

Movie posters are nearly identical to book covers concerning their aim and the content that they present. Classification has been done based considering only the poster, as well as poster along with movie title, plot and other metadata in [11]-[14]. It was also attempted to predict the box office collections of a movie based on movie poster and metadata in [13]. A similar analysis was performed on book covers as to whether they could be used a determiner of book sales and popularity in [29], and whether a Generative Adversarial Network (GAN) could be trained to simulate popular book covers. Genre classification for artistic style has also been explored in the field of art in [15]. For text modality, the same authors as [8] applied Doc2Vec and Multinomial Naïve Bayes to textual descriptions of books in [26]. Across all the considered genres, Doc2Vec achieved the highest accuracy. Multi-label binarizer was used for multi-genre prediction based on book plot summaries in [28]. Considering only book titles in [9], a variety of algorithms were evaluated, with LSTM giving the best accuracy by a very narrow margin. E-commerce product title classification in [30] used SVM and also revealed the harmful effects of applying typical conventional text classification procedures on short text. We looked at various neural network approaches to text classification in [29], where different CNNs are trained on pre-trained word vectors for sentence-level classification tasks. Many approaches rely on text as a single modality as most book metadata is represented as textual features. This approach often shows good results and yields low inference time, as processing text typically is faster than images. By including the information encoded in the images, we can expect the classification performance to improve. There have been a variety of approaches to image-text multimodal classification. Encoded text is embedded onto an image to obtain an information enriched image in [31]. The features of resulting images are learned by a CNN for classification purpose. Labels associated with images are used for semi-supervised learning in [32], where a Multiple Kernel Learning (MKL) classifier was first learned on both image content and keywords associated with the image, after which SVM is trained on visual features from MKL output values. A fusion process is introduced in [21] to combine the outputs of image and text models. Here a regularization parameter, λ, was used to control the balance between the two sub-models. To integrate text and image cues at feature level, [22] used an attention mechanism where image feature is used to attend over the text feature, and this transformed text feature is fused with image features for classification.

## III. DATASET

We used the publicly available dataset of book covers provided by [2]. The raw dataset contains information regarding the book cover title, authors, main category, and link to the images of 57,000 book covers equally divided among 30 main categories from Amazon.com. In cases where the book cover was assigned to multiple categories, one category has been randomly selected.

The book cover dataset is quite complex when we compare it to standard image classification datasets like MS-COCO [35] and ImageNet [36]. It is quite difficult even for humans to classify. Most of the categories are distinguishable based on the textual description, while some categories contain significant visual cues such as recognizable objects for its discrimination [19].

## IV. BOOK GENRE CLASSIFICATION

### A. Proposed Pipeline

In this paper we propose two neural network models, one for early fusion and one for late fusion. The proposed network structures are described in Fig. 2 and Fig. 3. They consist mainly of 3 parts: Text extraction from image, text and visual feature extraction, and multimodal classification based on either feature fusion or decision fusion. The first part aims at detecting and recognizing the text present on the cover. The second part transforms each input source into a feature representation. The final part combines the dual-modality representations into a single representation. The main components of the proposed method are described below.

### B. Text Detection

To extract the text present on the book covers different methods were tested. For all methods, the images were first converted to grayscale as colour information is not required here. Tesseract OCR with pytesseract wrapper was used for text detection and recognition. First, we simply applied some basic erosion and dilation before text detection, however the text detected was too noisy. We also applied edge detection methods as in [27]. Since the cover text is similar to scene text in terms of noise, the EAST model introduced in [37] was also tried with not dissimilar results. Finally, we chose to apply 4 different binary thresholds to the grayscale image and combined the text detected from each new image generated, as shown in Fig. 1.
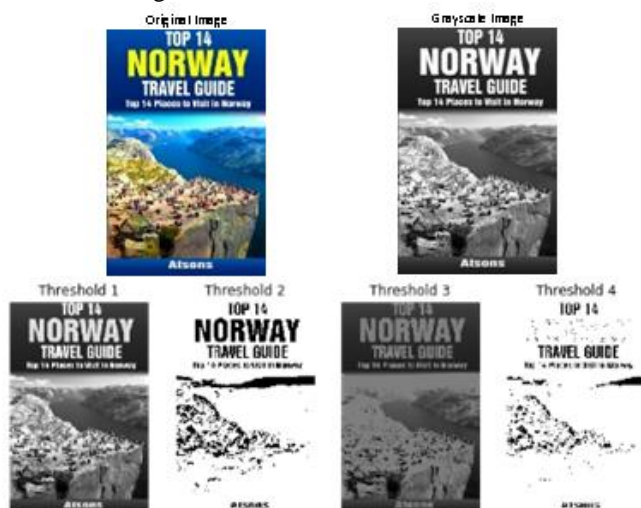


**Fig. 1. Sample book cover with the application of different thresholds**

### C. Text Recognition

The application of thresholds was modified to trying the fourth threshold only if no text was detected in any of the first three thresholds. This proved to be more informative, however, it suffered from the problem of overlapping – since basically the same image is used to detect text multiple times, the same text is detected repeatedly. This is easily remedied, however sometimes the application of threshold results in warping of a letter or two, resulting in the detection of a seemingly different word which results in noisy text.

### D. Text Preprocessing

Since our text data is of short length, with the longest title not having more than 20 words, as stated in [38], typical text preprocessing methods such as stopword removal, lemmatization and stemming do more harm than good. Instead, we opted to simply remove as much noise as possible by removing all punctuation and retaining words at least 3 characters long.

### E. Text Representation

CountVectorizer along with TfidfTransformer was used to obtain feature vectors for our text. We also looked at Keras' Tokenizer API as well as pre-trained GloVe word embeddings to create an embedding matrix that is used as input for text classification CNN.

### F. Text Classifier

Naïve Bayes classifiers are a family of classifiers that are based on the popular Bayes' probability theorem. They are known for creating simple yet well-performing models, especially in the fields of text classification. We use Multinomial Naïve Bayes classifier with alpha value 0.39 to get the best accuracy on text classification. We also attempted to design a custom CNN, where the word embeddings were given as input to an Embedding Layer. The output of this layer was given to separate convolutional layers with 256 filters and kernel sizes 3, 4 respectively. The output of each of these layers is passed through a maxpooling layer, concatenated and then flattened since we wish to give it as input to a dense layer. The dense layers contain 256 units. After this, we added a dropout layer with rate 0.5, and finally a 30 unit dense layer with softmax activation.

### G. Image Preprocessing

The cover images were resized to 299x299 since pre-trained weights for those dimensions are available for Inception-v3 CNN. We used the same as the input to our custom CNN. Some data augmentation such as rescaling was also done to increase the number of available training and testing images.

### H. Image Classifier

Inception CNN, developed by researchers at Google, has obtained state-of-the-art results in the field of large-scale multiclass image classification. It has also given excellent results for transfer learning. We used the base layers of Inception-v3 [6] to extract features from the cover images. With this we tried 2 methods: remove the bottleneck and replace it with a dense layer, followed by a Dropout layer and finally a dense layer with softmax activation to obtain a prediction for the 30 classes. We also took the 2048 features obtained from the last average pool layer and tried to reduce its dimensionality. Upon applying Principal Component Analysis (PCA), we reduced the number of features to 229. These features were then used to train an SVM classifier. We performed Grid Search to determine the best set of parameters for the classifier. We also designed our own CNN,

where the input image is in parallel passed to both a 2D convolutional layer with 256 filters and a max pooling layer of pool size (2, 2). The output of each is again given to convolutional layers after which the features obtained in each limb are concatenated.

These features are flattened and then passed to a dense layer with softmax activation to get the final output.

### I. *Early Fusion*

As mentioned previously, in early fusion the different modes are combined at the feature level and the combined features are used to train the classifier. We take the flattened vectors obtained in intermediate stages from our image and text CNNs. These vectors are concatenated to give our multimodal feature vectors. The vectors then go through dense layers with 128 and 64 units respectively. Finally, after a dropout layer with rate 0.5, we have a dense layer with softmax activation to give our predictions.

### J. *Late Fusion*

We opted to use the predictions generated by the combined Inception + SVM model for image classifier predictions and the predictions of the Naïve Bayes model for text predictions. We tested different policies to generate the final genre prediction, such as adding the class-wise probabilities, product max and argmax policies. We finally chose product-max heuristic policy to give the final prediction. This is equivalent to:

$$P_{\text{product-max}} = \arg\max (o^{\text{image}} \odot o^{\text{text}})$$
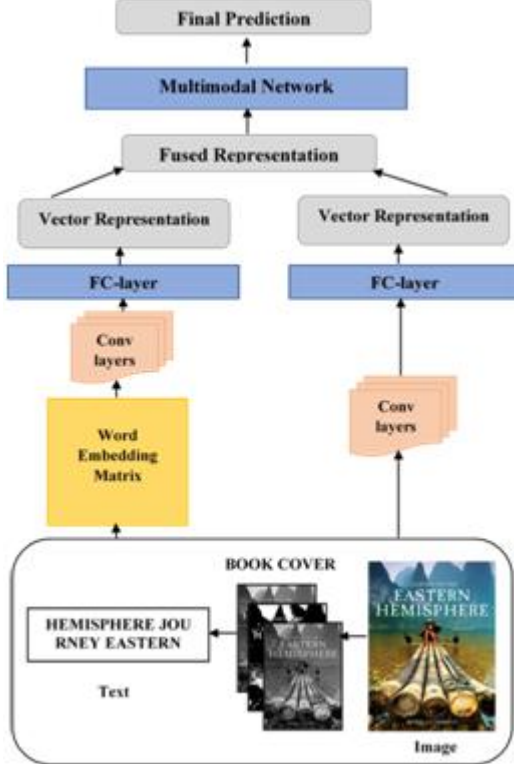
Here P refers to the final prediction probability.
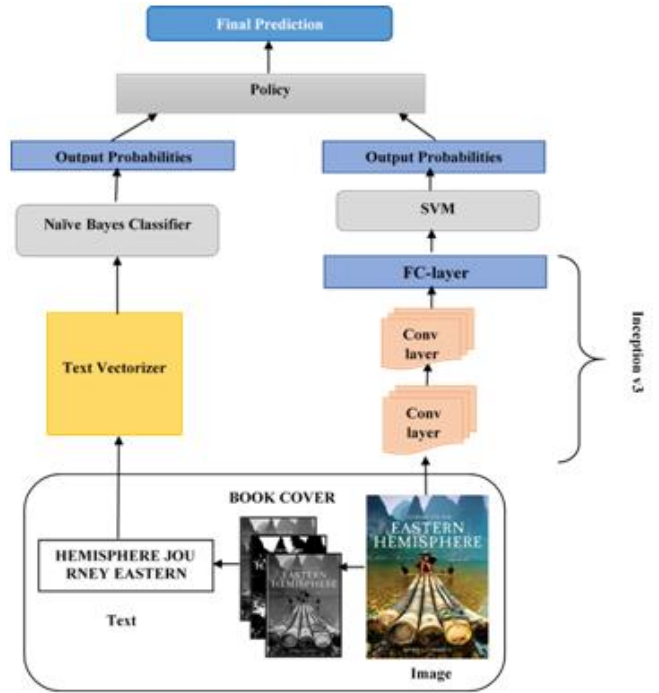


Fig. 2. Early Fusion Architecture



**Fig. 3. Late Fusion Architecture**

## V. RESULTS AND DISCUSSION

### A. *Image only*

Using only Inception CNN with some hyperparameter tuning we managed to achieve a Top 1 accuracy of 30%. Extracting the features from Inception and training an SVM classifier on those features improved the accuracy by a few percent. For a 30 class image classifier, the results are not exceptional. This might be due to dissimilarity of covers belonging to the same genre which causes the model to overfit on the training data. The categories such as *Cooking, Food and Wine*, *Sports and Outdoors* have some kind of common imagery across most of the books belonging to the category, such as food and drinks in the first case and sports equipment/sportspersons in the latter. However, categories such as *History* or *Self Help* do not have any specific distinguishing features which make them harder to classify.

### B. *Text only*

The best accuracy using only text is 55.4%. Here also overfitting might have occurred. Considering the complete text vocabulary, only a very small subset of it can act as keywords denoting a particular genre. Especially in fiction genres, titles can be very ambiguous. This probably means we must look at how better to capture semantic meaning and relation between different words present in a single book's cover text to get better classification results.

### C. *Image and Text*

#### 1) *Early Fusion*

Early fusion did give an improvement over classification based only on image, however it was not more accurate than text-only classification. It is evident that the features being extracted are not very well suited

**Table- I: Experimental Results from Tested Machine Learning Algorithms**

| Modality | Method | Accuracy |
|---|---|---|
| Image only | InceptionV3 | 30.1% |
| | InceptionV3 + SVM | 34.6% |
| Text only | Naïve Bayes | 55.4% |
| | CNN | 53.1% |
| Image + Text | Feature fusion (early fusion) | 56.2% |
| | Decision fusion (late fusion) | 60.1% |

to this task. Since we used custom CNNs to extract the features of both image and text, we must reexamine our network architecture to further design a neural network that better extracts the features from both images and text for book covers.

*2) Late Fusion*

The multimodal late fusion model gave us the overall best accuracy. Typically late fusion is said to be optimal in case of unrelated modalities, but in our case perhaps late fusion served to capture different features from image and text that did correspond to the genre it belonged to. Product-max heuristic policy gave us the best accuracy. The class-wise probabilities obtained were multiplied element-wise to give the final probability. This showed that late fusion models, despite being simple, worked quite well for our classification task.

## VI. CONCLUSION

In this paper, we have trained models to predict the genre of a book based on visual and textual modalities of the cover and hence proved that it is possible to draw a relationship between book cover images and genre using automatic recognition. Our experimental results have demonstrated that textual features of books are much more informative for categorization when compared to visual features, and that combining shows a marked improvement in accuracy. The cover design process can benefit from this research as analysis of correctly and incorrectly classified covers reveal some underlying information as to what visual features help to identify books belonging to a particular genre. It could also be used to help in e-commerce platforms by providing automatic genre suggestions. Since the model is not completely accurate, the automatic suggestion can be displayed as a prompt, with an option for the user to modify it. It may also be helpful for recommendation services and in computer vision fields.

Future research will be put into further analysis of the characteristics of the classifications and the features determined by the CNN network in an attempt to design a network that is optimized for this task. We can try to fine-tune hyperparameters to improve accuracy. Also, it can be seen that the image and text features themselves are mostly independent of each other. This can be improved by using some sort of attention mechanism – either attending over the image vectors based on the text or vice versa. This may combine the modalities more effectively.

## REFERENCES

1. H. Chiang, Y. Ge, and C. Wu, "Classification of book genres by cover and title," 2016.
2. B.K. Iwana, S.T.R. Rizvi, S. Ahmed, A. Dengel, and S. Uchida, 2016. Judging a Book by its Cover. arXiv preprint arXiv:1610.09204.
3. Y. LeCun, L. Bottou, Y. Bengio and P. Haffner, 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), pp.2278-2324.
4. M. D. Zeiler and R. Fergus Zeiler. "Visualizing and understanding convolutional networks." *European conference on computer vision*. Springer, Cham, 2014.
5. D. Ciresan, U. Meier, and J. Schmidhuber. "Multi-column deep neural networks for image classification." *2012 IEEE conference on computer vision and pattern recognition*, pp. 3642-3649. IEEE, 2012.
6. C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in Proc. IEEE Conf. Comp. Vision and Pattern Recognition, 2015, pp. 1–9.
7. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
8. P. Buczkowski, A. Sobkowicz, and M. Kozlowski. "Deep Learning Approaches towards Book Covers Classification." ICPRAM. 2018.
9. E. Ozsarfati, E. Sahin, C. J. Saul, and A. Yilmaz, 2019, February. Book Genre Classification Based on Titles with Comparative Machine Learning Algorithms. In *2019 IEEE 4th International Conference on Computer and Communication Systems (ICCCS)* (pp. 14-20). IEEE.
10. Y. Shinahara, T. Karamatsu, D. Harada, K. Yamaguchi, and S. Uchida, 2019. Serif or Sans: Visual Font Analytics on Book Covers and Online Advertisements. *arXiv preprint arXiv:1906.10269*.
11. S. Sirattanajakarin, and P. Thusaranon. "Movie Genre in Multi-label Classification Using Semantic Extraction from Only Movie Poster." ACM, 2019.
12. G. Barney, and K. Kris, "Predicting Genre from Movie Posters"
13. M. Shafaei, A. P. López-Monroy, and T. Solorio. "Exploiting Textual, Visual and Product Features for Predicting the Likeability of Movies." (2019).
14. W. Chu and H. Guo. "Movie genre classification based on poster images with deep neural networks." Proceedings of the Workshop on Multimodal Understanding of Social, Affective and Subjective Attributes. ACM, 2017.
15. A. Karnewar, A. Kanawaday, C. Sawant, Y. Gupta, Classification of Abstract Images Using Machine Learning, 2017
16. C. Rigaud, S. Pal, J. Burie and J. M. Ogier, Toward speech text recognition for comic books, 2016
17. F. Rayar and S. Uchida. "Comic Text Detection Using Neural Network Approach." International Conference on Multimedia Modeling. Springer, Cham, 2019.
18. Y. Daiku, M. Iwata, O. Augereau, and K. Kise. "Comics Story Representation System Based on Genre." 2018 13th IAPR International Workshop on Document Analysis Systems (DAS). IEEE, 2018.
19. S. Jolly, B. K. Iwana, R. Kuroki and S. Uchida, "How do Convolutional Neural Networks Learn Design?," 2018 24th International Conference on Pattern Recognition (ICPR), Beijing, 2018, pp. 1085-1090, doi: 10.1109/ICPR.2018.8545624.
20. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
21. L. Åberg. "Multimodal Classification of Second-Hand E-Commerce Ads." 2018.
22. X. Bai, M. Yang, P. Lyu, Y. Xu and J. Luo, "Integrating Scene Text and Visual Appearance for Fine-Grained Image Classification," in IEEE Access, vol. 6, pp. 66322-66335, 2018, doi: 10.1109/ACCESS.2018.2878899.
23. F. Zhu, Z. Ma, X. Li, G. Chen, J.T. Chien, J.H. Xue, J.H. and J. Guo. "Image-text dual neural network with decision strategy for small-sample image classification." Neurocomputing 328 (2019): 182-188.
24. Y. Kim. "Convolutional neural networks for sentence classification." arXiv preprint arXiv:1408.5882 (2014).

534

25. B. C. S. Berno, A. C. Gabardo, L. T. Hattori, M. Gutoski, A. de Souza Inácio and H. S. Lopes, "A Framework for Analyzing Book Covers and Co-purchases using Object Detection and Data Mining Methods," 2019 IEEE Latin American Conference on Computational Intelligence (LA-CCI), Guayaquil, Ecuador, 2019, pp. 1-6, doi: 10.1109/LA-CCI47412.2019.9037060.
26. A. Sobkowicz, M. Kozłowski, and P. Buczkowski. "Reading Book by the Cover—Book Genre Detection Using Short Descriptions." *International Conference on Man–Machine Interactions.* Springer, Cham, 2017.
27. S. K. Nagar. "Extracting the Book and Writer's name from the cover image of the book(English) using image processing methods," unpublished.
28. S. Bali, and A. Pramod. "Book Genre Classification", unpublished.
29. Y. Xu, W. Boyang, and J. Cao. "Judging a Book by its Cover: A Modern Approach," unpublished.
30. H.F., Ho, C.H., Arunachalam, P., Somaiya, M. and Lin, C.J. "Product title classification versus text classification." Csie. Ntu. Edu., 2012, Tw, pp.1-25.
31. I. Gallo, A. Calefati, S. Nawaz and M. K. Janjua. "Image and Encoded Text Fusion for Multi-Modal Classification." *2018 Digital Image Computing:Techniques and Applications (DICTA),* IEEE. 2018.
32. Guillaumin, M., Verbeek, J. and Schmid, C., 2010, June. Multimodal semi-supervised learning for image classification. In 2010 IEEE Computer society conference on computer vision and pattern recognition (pp. 902-909). IEEE.
33. Q. Ye and D. Doermann, "Text detection and recognition in imagery: A survey," IEEE Trans. on Pattern Anal. and Mach. Intell., vol. 37, no. 3, pp. 1480–1500, 2015.
34. Y. Zhu, C. Yao, and X. Bai, "Scene text detection and recognition: Recent advances and future trends," Frontier of Computer Science, vol. 10, no. 1, pp. 19–36, 2016.
35. T.Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár and C.L. Zitnick. Microsoft coco: common objects in context. In: European conference on computer vision. Springer; 2014. p. 740–55.
36. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei.. ImageNet large scale visual recognition challenge. Int J Comput Vis (IJCV). 2015;115(3):211–52. https://doi.org/10.1007/s11263-015-0816-y.
37. X. Zhou, C. Yao, H. Wen, Y. Wang, S. Zhou, W. He and J. Liang. "EAST: an efficient and accurate scene text detector," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.* 2017.

**Bhanushree K J** is working as Assistant Professor in the Department of Computer science and Engineering at Bangalore Institute of Technology. Her areas of interest are Image processing and Machine Learning. She is currently pursuing research in the area of Face Recognition.

## AUTHORS PROFILE

**Rajasree Jayaram** is pursuing her Bachelor's Degree in Computer Science and Engineering from Bangalore Institute of Technology (VTU), Bengaluru, Karnataka. Her areas of interest include Neural Networks and Natural Language Processing.

**Harshitha Mallappa** is pursuing her Bachelor's Degree in Computer Science and Engineering from Bangalore Institute of Technology (VTU), Bengaluru, Karnataka. Her areas of interest include Machine Learning and Neural Networks.

**Pavithra S** is pursuing her Bachelor's Degree in Computer Science and Engineering from Bangalore Institute of Technology (VTU), Bengaluru, Karnataka. Her areas of interest include Neural Networks and Natural Language Processing.

**Munshira Noor B** is pursuing her Bachelor's Degree in Computer Science and Engineering from Bangalore Institute of Technology (VTU), Bengaluru, Karnataka. Her areas of interest include Machine Learning and Neural Networks.

535