

# How do Convolutional Neural Networks Learn Design?

Shailza Jolly\*, Brian Kenji Iwana†, Ryohei Kuroki†, Seiichi Uchida†

\*University of Kaiserslautern, Kaiserslautern, Germany

Email: sjolly@rhrk.uni-kl.de

†Department of Advanced Information Technology, Kyushu University, Fukuoka, Japan

Email: {brian, kuroki, uchida}@human.ait.kyushu-u.ac.jp

**Abstract**—In this paper, we aim to understand the design principles in book cover images which are carefully crafted by experts. Book covers are designed in a unique way, specific to genres which convey important information to their readers. By using Convolutional Neural Networks (CNN) to predict book genres from cover images, visual cues which distinguish genres can be highlighted and analyzed. In order to understand these visual clues contributing towards the decision of a genre, we present the application of Layer-wise Relevance Propagation (LRP) on the book cover image classification results. We use LRP to explain the pixel-wise contributions of book cover design and highlight the design elements contributing towards particular genres. In addition, with the use of state-of-the-art object and text detection methods, insights about genre-specific book cover designs are discovered.

## I. INTRODUCTION

Visual design renders specific impressions to transmit information which enriches the product’s value. However, these visual designs despite of being important are not analyzed objectively or statistically. Analyzing these visual designs enables us to understand the contained information carried by them.

An interesting target of visual design analysis is book cover image design where the design of a book cover can infer the genre. Each book cover is carefully designed by typographers and their designs represent the book contents in an intuitive way for better sales. This association of books to specific genres is based on the differences in their underlying book cover designs [1]. The slight change in book cover design can reflect changes in book genre which makes design learning a challenging task for book covers.

In order to understand the design elements used for machine aided book cover classification, we employ Convolutional Neural Networks (CNN) [2]. In recent years, CNNs have achieved state-of-the-art results in isolated character recognition [3], [4] and large-scale image recognition [5], [6]. Notably, Iwana et al. [1] demonstrated that CNNs can be used for genre classification based on book cover image, although with a high level of difficulty. However, that study was subjective and not enough explanation is given as to why the CNN performed as it did.

To interpret the reasoning behind a CNN’s prediction we used a method called Layer-wise Relevance Propagation (LRP)

[7]. LRP decomposes output function on its input variables and highlights input pixels contributing towards the network decision. It produces a layer-wise relevance heatmap by recursively multiplying the relevance of higher layers by the normalized feature maps of the target layer. The heatmaps can help us to discover the input image elements which have an effect on the classification result.

The main contributions of this paper are threefold. Firstly, we classified the book cover images using one-vs-others classification with CNNs. Secondly, the models built by the CNNs are analyzed using LRP. With LRP, we demonstrate design elements specifically relevant to classification of the book cover images. We show that certain objects have a strong relevance to particular genres. Finally, we use state-of-the-art object detection and text detection methods, namely Single Shot Multibox Detector (SSD) [8] and Efficient and Accurate Scene Text Detector [9], to quantitatively enforce the results found by LRP. This reveals the specific elements in which CNNs classify book cover images for genre classification.

The organization is as follows. Section II provides related works in design understanding and genre classification as well as feature visualization of CNNs. Section III reviews the data and tools used for understanding book cover design. Section IV presents analysis of CNN’s understanding of book cover design. In Section V, we demonstrate the use of LRP combined with SSD and EAST for quantitative analysis. Finally, Section VI draws a conclusion.

## II. RELATED WORK

### A. Genre Classification

Artistic style understanding and subjective genre classification is a budding field in machine learning. For example, recent attempts have been done to identify artistic styles and quality of paintings and photographs [10], [11] with neural network models. In addition, there have been trials to classify music by genre [12], [13], book covers by genre [1], movie posters by genre [14], paintings by genre [15], and text by genre [16], [17]. Also, in a general sense, document classification can be considered genre classification and deep CNNs are the state-of-the-art in the document classification domain [18]–[20].

## B. Visualization inside of CNNs

There is a desire to visualize features and determine pixel-wise attention and relevance within the hidden layers of CNNs. However, this is not a straightforward task [21]. Erhan et al. [21] proposed using gradient descent to maximize a node's activation to visualize the employed features. Similar work has been done for large-scale image classification [22]. Zeiler and Fergus [23] used deconvolutional neural networks to visualize features learned by CNNs. In addition, they created heatmaps by monitoring class changes systematic cover up of portions of the images. Class Activation Maps (CAM) [24], GradCAM [25], and GradCAM++ [26] reveal the parts of images which are most important to a class using global average pooling (GAP).

Recently, LRP has been used in the fields of text [27] where classification scores were projected back to input features for extracting relevant words for a specific prediction. The method has also shown successes in model understanding in fields of sentiment analysis [28], action recognition [29], and age and gender classification [30]. As far as the authors are aware, this is the first time LRP has been used for the understanding of genre or design classification.

## III. DATA AND TOOLS FOR UNDERSTANDING BOOK COVER DESIGN

### A. Amazon Book Cover Dataset

We used the *Book Cover Image to Genre* dataset<sup>1</sup> Task 1.A. The dataset consists of 57,000 book cover images divided into 30 classes of equal sizes. In the experiments, we used the predefined training set and test set modified for one-vs-others classification. In this way, genre-wise training sets were prepared with an equal distribution of positive and negative data samples.

### B. Convolution Neural Networks

CNNs are able to tackle image recognition by implementing convolutions of learned filter-like shared weights which maintain the structural qualities of images while acting as feature extractors [2]. For the experiment, we implement CNNs to tackle book genre classification. To use the book cover images with a CNN, they were preprocessed by scaling them to  $112 \times 112$  pixels by 3 color channel images and by normalizing the values to be between -1 and 1. The CNN used for the experiments has six convolutional layers with Rectified Linear Units (ReLU) activations and a softmax output layer. The convolutional layers consisted of three layers of 10 nodes with  $5 \times 5$  convolutional filters, one layer of 25 nodes with a  $4 \times 4$  filter, one layer of 50 nodes with a  $3 \times 3$  filter, and one layer of 100 nodes with a  $1 \times 1$  filter. A  $2 \times 2$  maxpooling layer with stride 2 was used between each convolution layer. Finally, the CNNs were trained using gradient descent with a batch size of 25 at a learning rate of 0.001 for 50,000 iterations.

The accuracy results for each genre is summarized in Fig. 1. In particular, the CNNs had difficulties with the

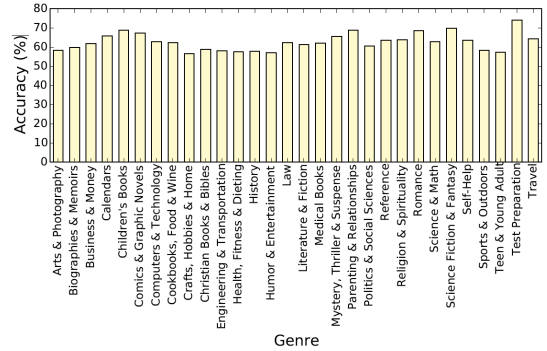


Fig. 1. CNN accuracy by genre.

reference classes, such as "Engineering & Transportation," "Health, Fitness and Dieting," "History," "Medical Books," and "Reference." Conversely, "Children Books," "Romance," and "Test Preparation" had high accuracies. However, more than just classification accuracy, the purpose of this paper is to understand why the CNN's performed as such and reveal the relevant parts of the images.

### C. Layer-wise Relevance Propagation

The LRP algorithm and the LRP toolbox [31] aims to explain the reasoning behind the decision made by a network model which allows its users to validate classifier results. LRP is mainly derived from Deep Taylor Decomposition [32], a method of decomposing network's output predictions onto its input variable. The results after such a decomposition is visualized in the form of a heatmap highlighting each pixel's importance for the prediction.

LRP explains output function, i.e. classifier's decision, which helps us to derive all of the crucial pixels for a particular prediction. In Fig. 2, the technique is shown in which the output value given by the network is decomposed backwards layer by layer until it reaches the input. This backward decomposition of network's prediction uses local redistribution rules for assigning relevance values  $R_i$  to each neuron contributing towards the output, namely

$$\sum_i R_i = \sum_j R_j = \dots = \sum_k R_k = f(x), \quad (1)$$

where  $f(x)$  is the prediction function,  $R_i$  is the relevance of node  $i$  in the target layer,  $R_j$  is the relevance of node  $j$  of the previous layer, and  $R_k$  is the relevance of node  $k$  of the highest layer. The total amount of relevance is conserved in this equation.

For the experiment, we used the  $\alpha - \beta$  decomposition formula defined by

$$R_i = \sum_j \left( \alpha \frac{(a_i w_{ij})^+}{\sum_i (a_i w_{ij})^+} + \beta \frac{(a_i w_{ij})^-}{\sum_i (a_i w_{ij})^-} \right) R_j, \quad (2)$$

where  $\alpha$  and  $\beta$  are hyperparameters to weight the positive values of  $\frac{(a_i w_{ij})^+}{\sum_i (a_i w_{ij})^+}$  and the negative values of  $\frac{(a_i w_{ij})^-}{\sum_i (a_i w_{ij})^-}$ , respectively. Furthermore,  $w_{ij}$  is the weight between nodes  $i$

<sup>1</sup><https://github.com/uchidalab/book-dataset>

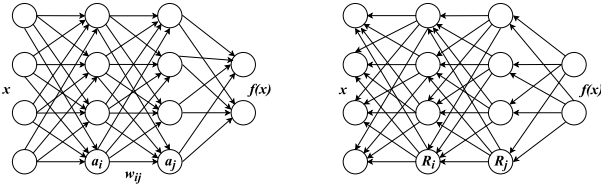


Fig. 2. Feed forward neural network with the (left) forward pass and the (right) backward relevance calculation. The function  $f(x)$  is the prediction outcome given input  $x$ . The variables  $a_i$  and  $a_j$  are the inputs for node  $i$  and  $j$ , respectively.  $R_i$  is the relevance of node  $i$  and  $R_j$  is the relevance of node  $j$ .

and  $j$  and  $a_i$  is the input to node  $i$ . This decomposition allows for the separation of the positive connections and the negative connections. Values inside positive bracket indicates propagation of activating input messages while negative weight connections indicate deactivating input values.

#### D. Single-Shot Multibox Detector

To develop a better understanding of the objects within book cover images, we employed SSD [8], a state-of-the-art deep neural network based object detection method. SSD is a feed forward CNN which produces a multi-scale collection of fixed size bounding boxes and scores for object detection within the boxes. A final non-maximal suppression step determines the final detections. The result of SSD is bounding box regions with object classification labels. Using SSD, it is possible to accurately detect multiple objects of different classes within images.

#### E. Efficient and Accurate Scene Text Detector

For humans, text is an important component of book covers; it is where the title, authors, and additional information is conveyed. However, a CNN may place a different importance on text than humans. Thus, to analyze the relevance of text in book covers, we use EAST [9] as a text detector. EAST uses a multi-channel Fully Convolutional Network (FCN) and non-maximal suppression on predicted geometric shapes to detect multi-orient text-line and word boxes.

### IV. HOW CNNs UNDERSTAND BOOK COVER DESIGN: QUALITATIVE ANALYSIS

In this section, we have presented LRP results from main genres. The analysis helped us to deduce book cover design elements contributing towards a prediction by CNN. We used  $\alpha - \beta$  decomposition formula with values of  $\alpha = 2$  and  $\beta = -1$  which is suggested for networks using ReLU activation functions because it emphasizes the positive elements and de-emphasizes the negative ones [7]. This is important due to the ReLU activation function setting negative values to zero. In the heatmaps generated by LRP under this decomposition, pixels adding positive contribution are represented in red color and the ones adding negative contribution are represented by blue color.

#### A. Sports & Outdoors

Under this genre, many book covers with pictures of players playing indoor and outdoor games were seen. Figure 3 (a) shows LRP results on these covers, which presents significance of player's picture on the cover. The first image in Fig. 3 (a) supports this fact with LRP being centered on players who are either playing a sport or showing some player like gesture, with car in background adding no contribution. The second image in Fig. 3 (a) emphasizes the animal's importance for this genre's prediction.

#### B. Engineering & Transportation

For this genre, almost all the covers with vehicle pictures on their covers were classified correctly by the network. With LRP in Fig. 3 (b), part of image containing cars or motorbikes seem to add more relevance than others. The last image in the Fig. 3 (b) presents the cases when contribution of person image was dominated by vehicle in the image.

#### C. Romance

Its obvious from the genre name that pictures of couples on the cover are going to have more relevance and LRP results showed this fact to be true. However, among pictures presented in Fig. 3 (c), LRP depicted girls to add more relevance than men or other things. The reason could reside in their physical appearance, hairs, and choice of dresses. The same was demonstrated in last picture of Fig. 3 (c) in which girl's hair are seen to add more relevance with zero relevance coming from animal part on book cover.

#### D. Children's Books

Almost all the children book covers contain pictures of cartoon characters. LRP on covers from this genre showed these cartoon characters to have higher relevance. An interesting result is shown in first picture of Fig. 3 (d), where person is depicted as an adversarial identity and importance of cartoons in cover is highlighted. Some covers showed more relevance for one object in the set of objects. Like, in Fig. 3 (d) some cartoons in last picture have higher relevance. It can be because of the object placement and their orientations. With the help of this information, one can make smart choices for different characters, cartoons and color patterns.

#### E. Cookbooks, Food & Wine Books

Book covers in this genre most commonly contained pictures of different kinds of food. The results in Fig. 3 (e) showed these food pictures as containment of higher relevance for this genre. However, carefully analyzing LRP results we discovered shapes of dishes like bowls or spoons adding significant relevance for the genre's prediction. So, this marks significance of dish shape designs on covers from this genre.

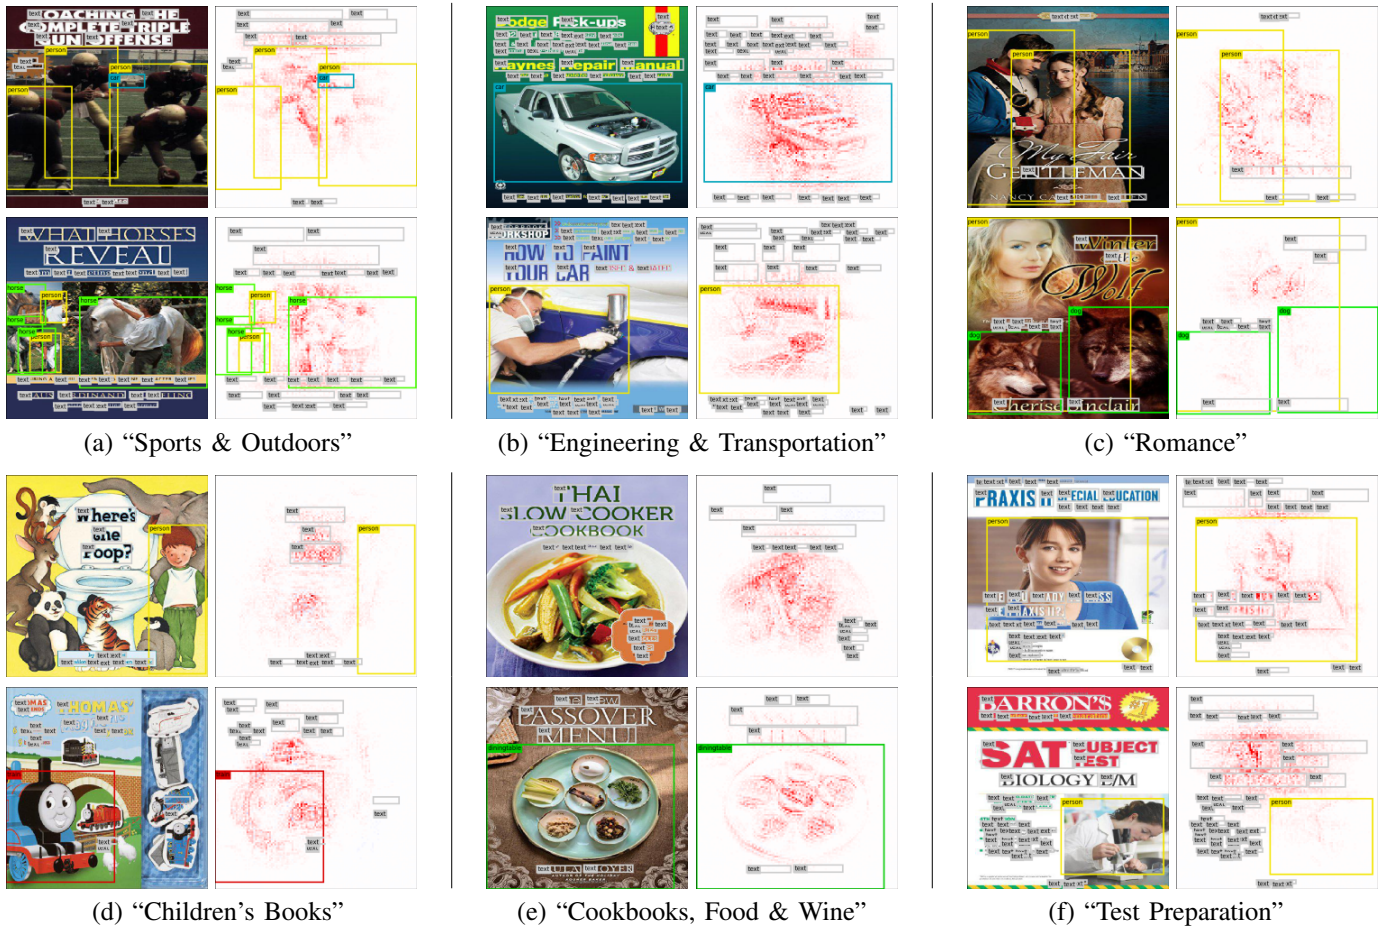


Fig. 3. Correctly recognized book covers. Object classes by SSD and text by EAST are highlighted.

#### F. Test Preparation

The genre contained covers with both text and pictorial information as shown in Fig. 3 (f). With most contribution coming from big text content on covers. Images of Fig. 3 (f) presents big texts to add more relevance than images of people. In first image of Fig. 3 (f), despite of big girl face, relevance is concentrated on text area of book cover.

Such analysis helped us to find design elements specific to the presented genres. To get more familiar with design, we also presented some cases where the network was not able to correctly classify the genre. Figure 4 shows some of these misclassifications, mainly from the presented genres. The correct genre names are written below the image. From the analysis presented above, one can easily decode the reason behind their misclassification because the designs on these book covers are not aligned with their genres which makes it obvious for network to mis-classify. Here, cover from "Sports & Outdoors" contains birds, "Romance" cover contains text, "Cookbooks, Food & Wine Books" contain no food picture and "Test Preparation" cover is also without any significant feature. LRP justifies all these covers misclassification by highlighting these mentioned objects contributing towards the "other" class in one-vs-others.

#### V. HOW CNNs UNDERSTAND BOOK COVER DESIGN: QUANTITATIVE ANALYSIS

##### A. Experiment Setup

In order to quantitatively analyze LRP, we propose using SSD to detect objects and EAST to detect text within the book cover images. We then use LRP to compare the relevance of objects bound by the detection methods. The SSD was trained on the 2012 PASCAL Visual Object Classes (VOC) Challenge dataset [33]. The VOC dataset contains 20 classes, including "person," six animal classes, eight vehicle classes, and seven indoor object classes. While SSD trained with VOC is intended for natural scene images, it can be used with book cover images because book covers often contain many of the shared classes, such as "person" and "car." Similarly, EAST was trained on the 2015 ICDAR Robust Reading Competition dataset [34] meant for scene text detection. Despite being trained for scene text, shown in Fig. 3, EAST performs remarkably well on book covers for detecting text.

To extract object and text bounding box information, the book covers were prepared by scaling the images to  $512 \times 512$  pixels by 3 color channels. It is important to note that the images used for SSD and EAST were larger than the images





Fig. 4. “Misclassified” book covers with correct genre names written below each book cover.

used by the CNN used for genre classification. This is due to the detection methods being much more effective at the higher resolution. To accommodate this, the bounding boxes were scaled post detection and projected onto the LRP heatmaps. The relevance of an object  $R_{obj}$  is calculated using the sum of the relevance within the bounding box, or

$$R_{obj} = \sum_{(n,m) \in \mathcal{B}} R(n,m), \quad (3)$$

where  $R(n,m)$  is the relevance at pixel coordinates  $(n,m)$  within bounding box  $\mathcal{B}$ .

### B. LRP with Object Detection

A macro view of the genres can be seen by viewing the average relevance of object classes. Figure 5 illustrates the average object-wise relevance of each object class as detected by SSD and EAST for each book genre using the test set book cover images. It should be noted that detected objects such as “bottle” and “tvmonitor” were overfit to certain book cover images because many books have plain covers which resemble bottle labels or televisions. However, this does not mean that the information is useless. For example, from Fig. 5, “bottle” is more relevant for reference and nonfiction genres where plain covers are common.

In addition, by examining the distribution of the  $R_{obj}$  of specific object classes, such as “person,” it is possible to create associations between genres and detected objects. For example, the relevance of “person”  $R_{person}$  for each genre is shown in Fig. 6. The figure demonstrates that detected “person”’s within certain genres are more relevant than other genres. For instance, the genres of “Romance” and “Mystery, Thriller & Suspense” put a high average relevance in “person.” This indicates that “person” is important for the CNNs of those categories. In addition, mentioned in Section IV-F and shown in Fig. 3 (f), people are common in “Test Preparation” but are not necessarily relevant. This is supported by Fig. 6 which indicates that on average, “person” has very little relevance. Distributions for the other object classes are provided in the supplemental material.

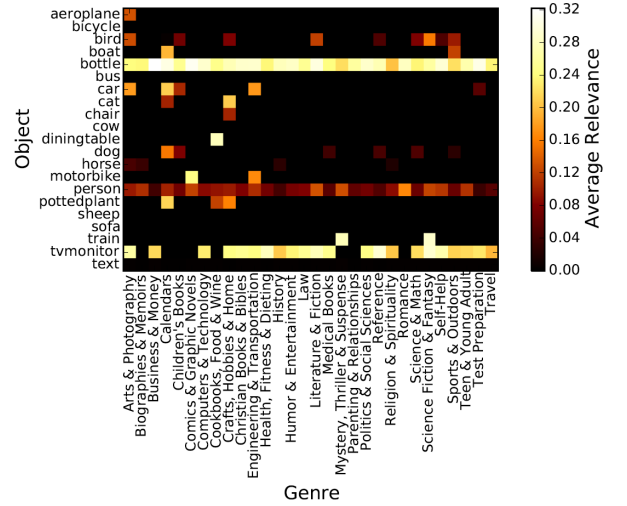


Fig. 5. Average object-wise relevance for text detected by EAST and each object class detected by SSD for each book genre. Only object-genre combinations with five or more data points are shown.

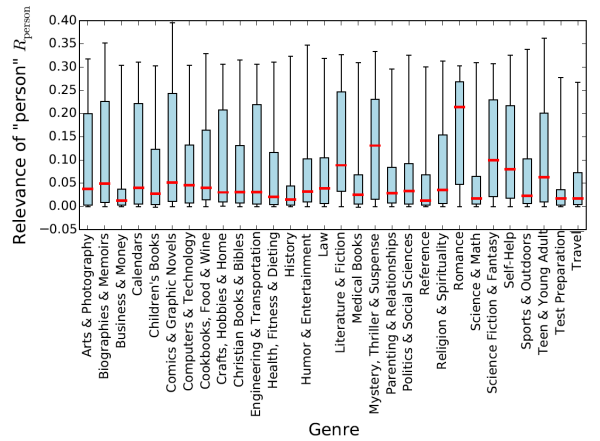


Fig. 6. Box plot of relevance of “person”  $R_{person}$  for each genre. The boxes represent the first through third quartile and the mean is in red. The whiskers mark the minimum and maximum datum.

### C. LRP with Text Detection

Figure 5 also reveals that the average relevance of text is low. The reasoning behind this phenomenon can be explained by Fig. 7. The figure shows that the majority of the detected text boxes have a very small relevance  $R_{text}$ , but there are some text boxes have a higher relevance. For most genres, the title text contains a significant amount of relevance determined by LRP, but the small descriptive text carries very little relevance. Figure 3 (f) in particular demonstrates this with the large title text having a high relevance and much of the smaller descriptive text having near zero relevance.

## VI. CONCLUSION

In this paper, we presented importance of design in book covers belonging to a specific genre. The application of LRP on the book cover dataset showed genre specific book cover features. The method described most relevant parts of input

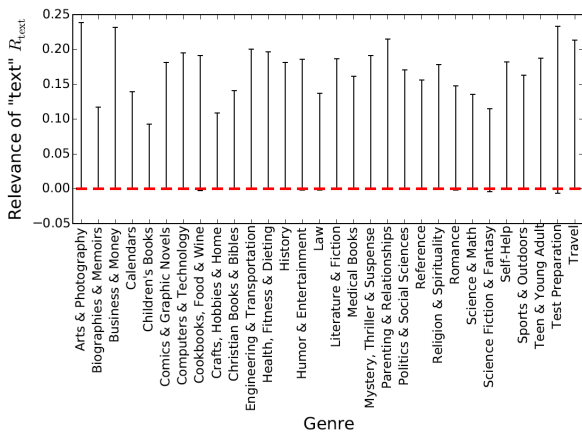


Fig. 7. Box plot of relevance of "text"  $R_{\text{text}}$  for each genre. The boxes represent the first through third quartile and the mean is in red. The whiskers mark the minimum and maximum datum.

book cover contributing towards a genre prediction by CNN. We also presented quantitative analysis of LRP using an object detection method, SSD, and a text detection method, EAST. The analysis further demonstrates that genre classification heavily relies on specific objects for each genres.

## VII. ACKNOWLEDGEMENT

This research was partially supported by MEXT-Japan (Grant No.J17H06100).

## REFERENCES

- [1] B. K. Iwana, S. T. R. Rizvi, S. Ahmed, A. Dengel, and S. Uchida, "Judging a book by its cover," *arXiv preprint arXiv:1610.09204*, 2016.
- [2] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [3] D. Ciregan, U. Meier, and J. Schmidhuber, "Multi-column deep neural networks for image classification," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*. IEEE, 2012, pp. 3642–3649.
- [4] S. Uchida, S. Ide, B. K. Iwana, and A. Zhu, "A further step to perfect accuracy by training cnn with larger data," in *Int. Conf. Frontiers in Handwriting Recognition*, 2016.
- [5] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [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. Comput. Vision and Pattern Recognition*, 2015, pp. 1–9.
- [7] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PLoS ONE*, vol. 10, no. 7, p. e0130140, 2015.
- [8] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European Conf. Comput. Vision*. Springer, 2016, pp. 21–37.
- [9] X. Zhou, C. Yao, H. Wen, Y. Wang, S. Zhou, W. He, and J. Liang, "East: An efficient and accurate scene text detector," in *IEEE Conf. Comput. Vision and Pattern Recognition*. IEEE, 2017, pp. 2642–2651.
- [10] R. Datta, D. Joshi, J. Li, and J. Z. Wang, "Image retrieval: Ideas, influences, and trends of the new age," *ACM Computing Surveys*, vol. 40, no. 2, p. 5, 2008.
- [11] S. Karayev, M. Trentacoste, H. Han, A. Agarwala, T. Darrell, A. Hertzmann, and H. Winnemoeller, "Recognizing image style," *arXiv preprint arXiv:1311.3715*, 2013.
- [12] G. Tzanetakis and P. Cook, "Musical genre classification of audio signals," *IEEE Trans. Speech and Audio Process.*, vol. 10, no. 5, pp. 293–302, 2002.

- [13] C. McKay and I. Fujinaga, "Automatic genre classification using large high-level musical feature sets," in *Int. Soc. of Music Inform. Retrieval*, vol. 2004, 2004, pp. 525–530.
- [14] W.-T. Chu and H.-J. Guo, "Movie genre classification based on poster images with deep neural networks," in *Proc. Workshop Multimodal Understanding of Social, Affective and Subjective Attributes*, 2017, pp. 39–45.
- [15] J. Zujovic, L. Gandy, S. Friedman, B. Pardo, and T. N. Pappas, "Classifying paintings by artistic genre: An analysis of features & classifiers," in *IEEE Int. Workshop Multimedia Sig. Process.* IEEE, 2009, pp. 1–5.
- [16] A. Finn and N. Kushmerick, "Learning to classify documents according to genre," *J. Amer. Soc. for Inform. Sci. and Technology*, vol. 57, no. 11, pp. 1506–1518, 2006.
- [17] P. Petrenz and B. Webber, "Stable classification of text genres," *Computational Linguistics*, vol. 37, no. 2, pp. 385–393, 2011.
- [18] L. Kang, J. Kumar, P. Ye, Y. Li, and D. Doermann, "Convolutional neural networks for document image classification," in *Int. Conf. Pattern Recognition*. IEEE, 2014, pp. 3168–3172.
- [19] A. W. Harley, A. Ufkes, and K. G. Derpanis, "Evaluation of deep convolutional nets for document image classification and retrieval," in *Int. Conf. Document Anal. and Recognition*, 2015, pp. 991–995.
- [20] M. Z. Afzal, S. Capobianco, M. I. Malik, S. Marinai, T. M. Breuel, A. Dengel, and M. Liwicki, "Deepdocclassifier: Document classification with deep convolutional neural network," in *Int. Conf. Document Anal. and Recognition*, 2015, pp. 1111–1115.
- [21] D. Erhan, Y. Bengio, A. Courville, and P. Vincent, "Visualizing higher-layer features of a deep network," *Tech. Rep. University of Montreal*, vol. 1341, p. 3, 2009.
- [22] C. Olah, A. Mordvintsev, and L. Schubert, "Feature visualization," *Distill*, 2017, <https://distill.pub/2017/feature-visualization>.
- [23] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European Conf. Comput. Vision*. Springer, 2014, pp. 818–833.
- [24] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, 2016, pp. 2921–2929.
- [25] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, "Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization," *arXiv preprint arXiv:1610.02391*, 2016.
- [26] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks," *arXiv preprint arXiv:1710.11063*, 2017.
- [27] L. Arras, F. Horn, G. Montavon, K.-R. Müller, and W. Samek, "What is relevant in a text document?": An interpretable machine learning approach," *PLoS ONE*, 2017.
- [28] L. Arras, G. Montavon, K.-R. Müller, and W. Samek, "Explaining recurrent neural network predictions in sentiment analysis," in *Proc. Workshop Computational Approaches to Subjectivity, Sentiment and Social Media Anal.* Association for Computational Linguistics, 2017, pp. 159–168.
- [29] V. Srinivasan, S. Lapuschkin, C. Hellge, K.-R. Müller, and W. Samek, "Interpretable human action recognition in compressed domain," in *IEEE Int. Conf. Acoustics, Speech and Sig. Process.*, 2017.
- [30] F. Arbabzadeh, G. Montavon, K.-R. Müller, and W. Samek, "Identifying individual facial expressions by deconstructing a neural network," in *German Conf. Pattern Recognition*, ser. Lecture Notes Comput. Science, B. Rosenhahn and B. Andres, Eds. Springer International Publishing, 2016, vol. 9796, pp. 344–354.
- [31] S. Lapuschkin, A. Binder, G. Montavon, K.-R. Müller, and W. Samek, "The lrp toolbox for artificial neural networks," *J. Mach. Learning Research*, vol. 17, no. 1, pp. 3938–3942, 2016.
- [32] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," *Pattern Recognition*, vol. 65, pp. 211–222, 2017.
- [33] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *Int. J. Comput. Vision*, vol. 88, no. 2, pp. 303–338, Jun. 2010.
- [34] D. Karatzas, L. Gomez-Bigorda, A. Nicolaou, S. Ghosh, A. Bagdanov, M. Iwamura, J. Matas, L. Neumann, V. R. Chandrasekhar, S. Lu et al., "Icdar 2015 competition on robust reading," in *Int. Conf. Document Anal. and Recognition*, 2015, pp. 1156–1160.