Original Article

Thai amulet recognition using vision transformers♣

Kittasil Silanon, Kullawat Chaowanawatee*, and Thitinan Kliangsuwan

*College of Computing, Prince of Songkla University,
Phuket Campus, Ka Thu, Phuket, 83120 Thailand*

Abstract

This study demonstrated the application of a deep learning method, namely the vision transformer model, to categorize Thai amulets from images. To construct a dataset, we selected famous Thai amulets that can be found easily online. Several vision transform models were used to train the system. We conducted five experiments to determine the most effective model. In Experiment 1, the model will receive direct training. Additional datasets will be produced in Experiment 2 using augmentation. The Test‑Time data Augmentation technique will be utilized in Experiment 3 to generate modifications in test dataset images. Finally, labeled and unlabeled data (Pseudo‑Labelling) were used simultaneously in each batch to train the model network in Experiment 4. The ensemble model combination is employed to improve model performance in Experiment 5. Furthermore, we developed a mobile application for capturing user‑provided image data. Users can upload images to the server for processing and access valuable information from the database, such as the backgrounds of amulets, creation ceremonies, and associated beliefs, to learn more about Thai amulets.

Keywords: Thai amulets, deep learning, vision transformers, image recognition

1. Introduction

Thai amulets, or "PHRA KHRUEANG," are symbols of religion that have been with the Thai people for a long time (Premjai & Matthew, 2022). Thais believe these amulets will protect them from danger, make them invincible, and bring good luck. Amulets are small Buddha images but can also feature images of monks, maestros, the Bodhisattvas, and other gods. Almost every Thai Buddhist has at least one amulet. It is expected that both young and older people will wear at least one amulet around their neck to feel closer to Buddha. In addition, amulets are considered a fusion of religion, art, and history, inspiring a significant market for collectors and scholars. However, the numerous types of amulets in Thailand are challenging for individuals to understand and recognize. Therefore, many research investigations on Thai amulets have been proposed. Chomtip,

Juti, Terapong, and Pimluk (2010) have introduced a system called the "Buddhist Amulet Recognition System" (BARS). Template matching is applied to recognize the Buddhist amulet image in the recognition process. The system precision is about 80% and takes 0.76 milliseconds per image for processing. Chomtip, Vachiravit, Pornpetch, and Nattida (2011) developed a system that can recognize Thai Buddhist sculptures. The system is called the Thai Buddhist Sculpture Recognition System (TBuSRS). In image recognition, the Euclidean distance is applied to recognize the Buddhist sculpture image. The precision rates of training and held out datasets were 90% and 72.38%, respectively. The average access time of the system is around 2.72 seconds per image. Chomtip and Natdani (2013) developed the Buddhist Amulet Coin Recognition System (BACRS). The BACRS applied a rule‑based technique and a genetic algorithm to recognize amulet coins. The precision rate of the system is 91.53%, and the average access time is 1.05 seconds per amulet image. Waranat and Chomptip (2014) used image processing and artificial neural networks (ANN) to recognize the digital images of Thai Buddhist amulets. The two basic ANN models employed in this research were perceptron and multi‑layer perceptron. The perceptron can recognize the amulet images correctly at 97.67%, and the multi‑layer perceptron can

---

*Corresponding author
 Email address: kullawat.c@phuket.psu.ac.th

produce the best classification result at 100%. Thanachai, Chalie, Toshiako, Pished and Kaneko (2014) demonstrated how to determine the kind of amulet from the taken image. Pre-processing techniques such as grayscale conversion, filtering, Prewitt edge detection, cropping, and resizing were applied to compute the image similarity values, and then the template with the highest similarity was identified as the recognition result. The system can identify at 60.6% accuracy. Weera and Thanasin (2015) proposed the development of Thai Buddha amulet identification using a simple local correlation feature. The feature is used with K-nearest neighbors for classification. The results showed that the proposed method gains a high recognition rate of about 89.35%. Narut and Sangthong (2020) applied the Convolution Neural Network (CNN) of deep learning to classify the Benjapakee Buddha amulet images. The CNN architecture is designed to be suitable for recognition. The model could correctly identify 80% of amulet images. Tanasai (2021) presented a case study of Thai amulet recognition using a geometric surface image. The geometric surface images and their color photographs were first used to train a Generative Adversarial Network (GAN). The trained generator model was then used to predict the geometric surface image from the input color image. The evaluation showed that the predicted geometric surface images contain less ambiguity than their color image counterparts under different lighting conditions. Chomtip and Varin (2022) used a pre-trained CNN model (ResNet50) to recognize Buddhist amulets. Furthermore, pre-trained CNN models ResNet18 and ResNet101 were trained to compare their accuracies. They also conducted cross-validation on held-out dataset, finding accuracy, sensitivity, specificity, and precision rates of 99%, 95%, 99%, and 95.41%, respectively.

Previous studies have made significant strides in amulet recognition using various technological methods. However, there still remains a gap in effectively categorizing these amulets using the latest advancements in deep learning. This paper seeks to fill this gap by demonstrating the application of Vision Transformers (ViTs), an advanced deep-learning model, in categorizing Thai amulet images. Unlike previous studies, which primarily focused on traditional image processing techniques and conventional neural network architectures, our research leverages the advanced capabilities of ViTs. We selected well-known Thai amulets to form a comprehensive dataset, ensuring the system's relevance and applicability. Our contribution is twofold. First, we explore the effectiveness of ViTs in the specific context of Thai amulet recognition, a domain yet to be extensively studied with such advanced technology. Second, we conduct five experiments designed to test different aspects of the ViTs model, including direct training, data augmentation, Test-Time Augmentation (TTA), Pseudo-Labelling, and ensemble learning. This comprehensive approach aims to optimize the model's performance in recognizing the diverse range of Thai amulets. Additionally, we developed a mobile application for capturing user-provided image data. Users can upload images to the server for processing and access valuable information from the database, such as the background of amulets, creation ceremonies, and associated beliefs, to learn more about Thai amulets. We organize the paper as follows. Section 2 describes the details of the dataset and ViTs model. We present the experiment details and results in Section 3. Finally, we conclude this paper in Section 4.

2. Dataset and Vision Transformers Model

In this section, we first describe the classes and information of the dataset. Then we describe the details of vision transformers (ViTs) model and its characteristic.

2.1 Thai amulet dataset

Since there are many different kinds of amulets, we selected the famous amulets of types "Phra Khun Phaen," "Phra Kring," "Phra Nang Phaya," "Phra Phong Suphan," "Phra Rod," "Phra Somdej," and "Phra Sum Kor" (Prowd, 2021). We collected Thai amulet images from public domain. The dataset was randomly partitioned into training and test sets, with 80:20 ratio, resulting in 1,631 and 404 images in the respective subsets. For our dataset, all samples were resized to match the input shape of the models under test. Consequently, they were normalized using ImageNet normalization statistics. Figure 1 shows images of selected Thai amulet classes.

Figure 1. Thai amulet types: (a) Phra Khun Phaen, (b) Phra Kring, (c) Phra Nang Phaya, (d) Phra Phong Suphan, (e) Phra Rod, (f) Phra Somdej, and (g) Phra Sum Kor

2.2 Vision transformers model

Vision transformers (ViTs) (Alexey *et al.*, 2021) are a type of deep learning neural network architecture that was originally developed for natural language processing (NLP) tasks. However, ViTs have also been shown to be effective for various computer vision tasks, such as image classification, object detection, and semantic segmentation. ViTs work by first converting images into a sequence of patches. Each patch is then embedded into a high-dimensional vector space. The embedded patches are then fed into a transformer encoder, which learns to represent the relationships between different parts of the image. The transformer encoder is a stack of self-attention layers. Self-attention is a mechanism that allows the model to learn long-range dependencies in the data. This is important for computer vision tasks, as it allows the model to learn relationships between different parts of the image, even if they are far apart. Once the transformer encoder has processed the embedded patches, the outputs are fed into a decoder, producing a vector representing the probability distribution of the image across the different classes. ViTs have several advantages over traditional CNN models for

computer vision tasks. Such as being more efficient in terms of both memory and computing requirements, better at learning long‑range dependencies in the data, more flexible, and easily adapted to different tasks. Currently, EVA‑02 (Yuxin *et al.*, 2023) is a powerful and versatile ViTs model that can be used for various applications. EVA‑02 has achieved state‑of‑the‑art results on many vision tasks while utilizing significantly fewer parameters and computing budgets than other vision transformers models. There are four variations of the EVA‑02 model, with sizes ranging from 6M to 304M parameters. These models perform magnificently in image recognition tasks (Table 1).

3. Experiments and Results

Five experiments were conducted to assess the classification performance of Thai amulet recognition. Essential parameters for all experiments are setup when training the model. The learning rate controls how quickly the model learns and is set to $2e^{-3}$. The batch size, the number of images the model processes in each iteration, is set to 16. The number of training epochs, the number of times the model sees the entire training dataset, is set to 30. The input images to train the model are resized to 224x224 pixels. Predictive model evaluation is done using k‑fold cross validation (Shanthababu, 2023). The dataset was divided into five folds. The model was then trained and evaluated five times, using a different fold as the validation set each time. This technique reduces the risk of overfitting and provides a more accurate estimate of the model's generalization performance.

3.1 Vision transformers training

In this experiment, we executed the EVA‑02 model in four distinct variations (Tiny, Small, Base, and Large). Each model was directly trained using the training dataset. The performance is evaluated using accuracy and weighted F1‑score. Accuracy is a simple measure of overall correctness. At the same time, the weighted F1‑score is a more sophisticated metric that considers class imbalances. Testing results are shown in Table 2. The large EVA‑02 model at the fifth fold achieved the best performance at 62.35% accuracy and 61.86% weighted F1‑score. As part of the class performance analytics, Figure 2 demonstrates that while the model has a reasonable recognition rate for some classes, it struggles with others. Notably, the "Phra Sum Kor" class has a high recognition rate, but "Phra Rod" and "Phra Somdej" appear to be frequently confused with "Phra Sum Kor". This suggests the model requires more distinguishing features to differentiate between amulets with similar design elements. Direct training without augmentation seems insufficient for the model to learn the necessary robust features. Figure 3 illustrates some errors of recognition. However, the recognition results could be more satisfactory because of the limited number of data records. As a result, in the following experiment, we will increase the number of training data to improve the model performance.

3.2 Data augmentation

Data augmentation is a technique used in deep learning to increase the size and diversity of a training dataset

Table 1. Four variations of EVA-02 model

EVA-02 model	Parameters (Million)	FLOPs (Billion)	Top-1 accuracy on image net (%)
EVA-02 Tiny	10M	0.6	76.2
EVA-02 Small	31M	1.6	77.5
EVA-02 Base	101M	3.3	79.0
EVA-02 Large	304M	7.6	81.4

Table 2. Model performance evaluation

Model	Fold	Test dataset	
		Accuracy	Weighted F1
Tiny	1	0.4877	0.4852
	2	0.4938	0.4896
	3	0.5185	0.5159
	4	0.5247	0.5158
	5	0.5432	0.5362
Small	1	0.4506	0.4405
	2	0.4352	0.4270
	3	0.4537	0.4501
	4	0.4599	0.4543
	5	0.5000	0.4981
Base	1	0.5556	0.5551
	2	0.5988	0.5896
	3	0.6049	0.5991
	4	0.5895	0.5858
	5	0.5833	0.5710
Large	1	0.5957	0.5811
	2	0.5247	0.5161
	3	0.6235	0.6076
	4	0.5957	0.5906
	5	0.6235	0.6186

artificially (Alexander *et al.*, 2020). This is done by creating new data from existing data using various transformations. In order to increase the size of training dataset, this experiment used horizontal/vertical flipping, space translation, random focus, noise addition, and slight rotation. Figure 4 illustrates data resulting from such augmentation. Testing results are shown in Table 3. The base EVA‑02 model at the fifth fold achieved the best performance at 66.05% accuracy and 65.35% weighted F1‑score. As part of the class performance analytics, Figure 5 demonstrates that the recognition rates for most classes and "Phra Rod" and "Phra Somdej" improved, indicating that augmentation helped the model generalize better. However, there still needs to be more clarity for some classes. These errors imply that while augmentation adds robustness, it may not address the model's sensitivity to intra‑class variability. Employing data augmentation increases the amount of data to train the model. However, we can increase the number of test data to improve model performance, as described in the following experiment.

3.3 Test-time data augmentation

Test‑Time data Augmentation (TTA) is a technique that involves applying random transformations to test images before making predictions (Masanari, 2021). This can improve the model's performance by making it more robust to

Figure 2.   Confusion matrix of the large EVA-02 at 5th fold

Figure 3.   Examples of labeling errors

Figure 4.   Example of data augmentation: (a) original, (b) scaled, (c) rotated, (d) translated, (e) noise, (f) contrast, and (g) defocused

variations in the input data. TTA is similar to data augmentation, typically used during training to increase the size and diversity of the training dataset. However, TTA is applied to test images, which are not used to train the model.

Table 3.   Model performance evaluation after data augmentation

Model	Fold	Test dataset	
		Accuracy	Weighted F1
Tiny	1	0.4846	0.4764
	2	0.4753	0.4655
	3	0.4907	0.4894
	4	0.5185	0.5173
	5	0.4784	0.4697
Small	1	0.4352	0.4324
	2	0.4475	0.4430
	3	0.5000	0.4948
	4	0.5031	0.4970
	5	0.5154	0.5100
Base	1	0.5494	0.5452
	2	0.6142	0.6061
	3	0.6142	0.6074
	4	0.6204	0.6090
	5	0.6605	0.6535
Large	1	0.6173	0.6118
	2	0.6574	0.6539
	3	0.5926	0.5764
	4	0.5679	0.5552
	5	0.5988	0.5961

After the test images have been augmented, the model makes predictions on each augmented image. The final prediction is then made by averaging the predictions from the augmented images. We compared the performance results of the test dataset without TTA from the previous experiment and the test dataset using the TTA technique of models. From Table 4, the base EVA-02 model at the fifth fold using the TTA technique achieved the best performance at 67.28% accuracy and 66.81% weighted F1-score. As part of the class performance analytics, Figure 6 demonstrates the classification results for all classes. The "Phra Somdej" class still improved in the results to an accuracy of 45%. Many

Figure 5. Confusion matrix of the base EVA-02 at 5th fold with data augmentation

Figure 6. Confusion matrix of the base EVA-02 at 5th fold with TTA technique

others had increased accuracy in this experiment, such as the "Phra Khun Phaen" class, "Phra Nang Phaya" class, "Phra Phong Suphan" class, and the "Phra Sum Kor" class, but issues remain in others like "Phra Rod". This suggests that TTA helps with generalization to some extent but may also introduce noise that can lead to misclassification in cases where the model is not well-tuned to the augmented data. In the following experiment, we use labeled and unlabeled data to improve the model performance.

3.4 Pseudo-labelling

Pseudo-labeling (Dong-Hyun, 2013) is a semi-supervised learning technique that can be used to improve the performance of models. It uses a model trained on a labeled dataset to predict labels for an unlabeled dataset. The predicted labels (pseudo-labels) are then used to retrain a model, which can achieve better performance than the model trained on the labeled dataset alone. In this experiment, the best performance model from the previous experiment was selected to create the pseudo-labels for unlabeled data. The labeled (training data) and pseudo-labeled data were used to retrain the EVA-02 model. Finally, the new retrained model called labels of the test dataset with the TTA technique. Table 5 shows the model performance after using pseudo-labeling. The large EVA-02 model at the fourth fold using the pseudo-labeling technique achieved the best performance at 69.80% accuracy and 69.12% weighted F1-score. As part of the class performance analytics, Figure 7 demonstrates the classification results for all classes. The classes "Phra Khun

Table 4.   Model performance evaluation after TTA technique

Model	Fold	Test dataset without TTA		Test dataset using TTA	
		Accuracy	Weighted F1	Accuracy	Weighted F1
Tiny	1	0.4846	0.4764	0.5525	0.5497
	2	0.4753	0.4655	0.5216	0.5119
	3	0.4907	0.4894	0.5710	0.5659
	4	0.5185	0.5173	0.5586	0.5560
	5	0.4784	0.4697	0.5185	0.5079
Small	1	0.4352	0.4324	0.4630	0.4524
	2	0.4475	0.4430	0.5216	0.5152
	3	0.5000	0.4948	0.5340	0.5297
	4	0.5031	0.4970	0.5617	0.5548
	5	0.5154	0.5100	0.5741	0.5687
Base	1	0.5494	0.5452	0.6296	0.6275
	2	0.6142	0.6061	0.6605	0.6499
	3	0.6142	0.6074	0.6667	0.6570
	4	0.6204	0.6090	0.6574	0.6502
	5	0.6605	0.6535	0.6728	0.6681
Large	1	0.6173	0.6118	0.6265	0.6130
	2	0.6574	0.6539	0.6543	0.6512
	3	0.5926	0.5764	0.6080	0.5957
	4	0.5679	0.5552	0.6204	0.6159
	5	0.5988	0.5961	0.5988	0.5986

Table 5.   Model performance evaluation after pseudo-labeling

Model	Fold	Pseudo-Labeling Test dataset using TTA	
		Accuracy	Weighted F1
Tiny	1	0.5421	0.5314
	2	0.5470	0.5297
	3	0.5891	0.5745
	4	0.5248	0.5147
	5	0.5569	0.5484
Small	1	0.5173	0.4997
	2	0.5272	0.5174
	3	0.5569	0.5446
	4	0.5891	0.5741
	5	0.5668	0.5555
Base	1	0.6559	0.6524
	2	0.6683	0.6582
	3	0.6634	0.6561
	4	0.6782	0.6703
	5	0.6782	0.6688
Large	1	0.6782	0.6701
	2	0.6634	0.6543
	3	0.6931	0.6850
	4	0.6980	0.6912
	5	0.6287	0.6152

Phaen," "Phra Kring," "Phra Phong Suphan," and "Phra Rod" improved. However, the confusion matrix indicates persistent misclassifications, for example, between "Phra Nang Phaya" and "Phra Phong Suphan". This could imply a need for better-quality pseudo-labels or a more refined approach to integrating unlabeled data into the training process. To enhance the recognition rate, in the final experiment, we will test how to use a combination of multiple models for better prediction results.

3.5 Ensemble learning

Ensemble learning (Jason, 2021) is a machine learning technique that combines the predictions from multiple models to produce a more accurate and robust prediction. It is based on the idea that a group of models can make better predictions than any individual model. The best model from the previous was selected as a base model (i.e., the large EVA-02 model at the fourth fold, trained using pseudo-labeling). We added one model to evaluate the recognition accuracy, and an averaging technique was utilized. Subsequently, the class with the maximum probability after averaging is selected as the final prediction. Table 6 demonstrates the recognition rate after ensemble learning. Combining the Large–Fold 4 and Large–Fold 3 models achieved the best performance, around 70.0%. As part of the class performance analytics, Figure 8 shows improved performance across most classes, suggesting that ensemble


Figure 7.   Confusion matrix of the large EVA-02 at 4th fold trained with pseudo-labeling

Table 6.   Performance evaluation of ensemble models

Ensemble learning	Accuracy	Weighted F1
Large – Fold 4		-
Add Large – Fold 3	0.7030	0.7070
Add Large –Fold3 and Large – Fold 1	0.6980	0.6914

methods can effectively reduce the impact of individual model biases. Nevertheless, some misclassifications persist, indicating that the constituent models in the ensemble may still share common weaknesses, such as overfitting certain features that need to be more indicative of the correct class. The "Phra Rod" and "Phra Somdej" classes appear to need additional data to improve the model accuracy further.

3.6 Mobile application for Thai amulet recognition

The application is named "Thai Amulet Recognition" (Figure 9 (a)). Users can upload and classify Thai amulet images with this application. We used the Flutter framework to build the frontend part of our mobile application, while Python and the Flask framework were employed to develop the backend part. The classification model is stored on the server. Users can use the application's main interface to capture Thai amulet images or choose from their mobile gallery (Figure 9 (b)). Subsequently, the application uploads the image to the server. The server then prompts the trained model to predict the most likely classes of Thai amulets. The output class retrieves additional details from the database, including images, class name, and background for learning more about Thai amulets (Figure 9 (c)).


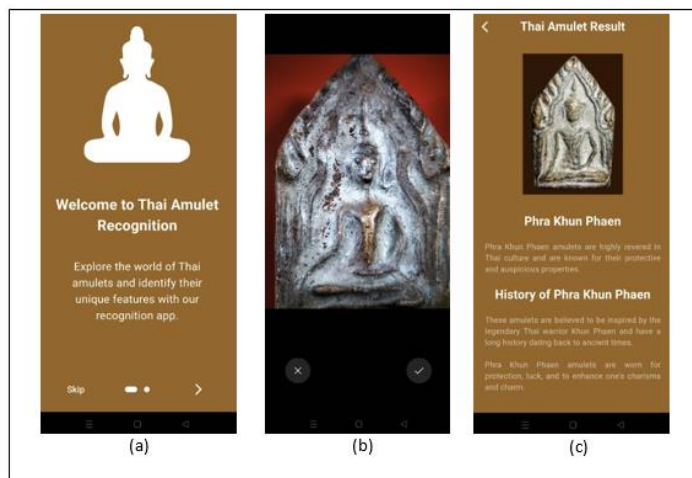Figure 8.   Confusion matrix of the ensemble model

Figure 9.   Mobile application for Thai amulet recognition: (a) main screen, (b) image upload, and (c) amulet information

4. Conclusions

In this research, we successfully developed and evaluated a series of model training experiments for Thai amulet recognition, employing a ViTs model alongside various techniques like data augmentation, TTA technique, pseudo-labeling, and ensemble learning. Our comprehensive approach demonstrated enhanced performance in amulet recognition and culminated in developing a user-friendly mobile application. This application allows users to upload amulet images and receive detailed information quickly. Our methodologies have the potential to be applied in broader contexts, such as cultural heritage preservation and educational tools, and even in enhancing the tourism industry by offering interactive and informative experiences related to Thai cultural artifacts. However, a limitation that makes our work less accurate is that the number of image data still needs to be more significant. Moreover, the current model's performance in varying environmental conditions, such as different lighting or angles, has yet to be extensively tested. Addressing these limitations will be crucial for enhancing the practicality and robustness of the recognition system. In future work, we will collect more data and create other experiments for several types of Thai amulets in order to enhance the recognition algorithm. We will compare the ViTs model used in this study and the CNN model. Another avenue for future research is the integration of user feedback mechanisms in our mobile application. This feature would allow users to provide real-time feedback on recognition accuracy, thereby enabling continuous learning and improvement of the algorithm.

Acknowledgements

References

Alexander B., Alex P., Eugene K., Vladimir I. I., & Alexandr A. K. (2020). Albumentations: Fast and flexible image augmentations. *Information*, *11*(2). doi:10.3390/info11020125

Alexey D., Lucas B., Alexander K., Dirk W. Xiaohua Z., Thomas U., . . . Neil Houlsby. (2021). An image is worth 16x16 words: transformers for image recognition at scale. Retrieved from https://doi.org/10.48550/arXiv.2010.11929

Chomtip, P., Varin, P. (2022). Buddhist amulet recognition by using RestNet50. *Srinakharinwirot Science Journal*, *38*(2).

Chomtip, P., Vachiravit, A., Pornpetch, I., & Nattida, P. (2011). Thai Buddhist sculpture recognition system (TBuSRS). *International Journal of Engineering and Technology*, *3*(4). 342-346 doi:10.7763/IJET.2011.V3.250

Dong-Hyun, L. (2013). Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. *Workshop on Challenges in Representation Learning, ICML, 3*(2), 896.

Iamkhorpung, P., & Kosuta, M. (2022). The relationship between Buddhist and animist amulets in contemporary Thailand: PHRA KHRUEANG and KHRUEANG RANG. *Kasetsart Journal of Social Sciences, 43*(1), 53–59. Retrieved from https://so04.tci-thaijo.org/index.php/kjss/article/view/256957

Jason, B. (2023, October 10). Ensemble learning. Retrieved from https://machinelearningmastery.com/tour-of-ensemble-learning-algorithms/

Kitiyanan, W., & Pornpanomchai, P. (2014). Thai Buddhist amulet recognition system. *Proceeding of the 2014 International Conference on Informatics and Advanced Computing, Bangkok* (pp. 9-14).

Masanari, M. (2012). Understanding test-time augmentation. *Proceeding of International Conference on Neural Information Processing* (pp. 558–569).

Narut, B., & Sangthong B. (2020). Classification of Benjapakee buddha amulets image by deep learning. *RMUTSB Academic Journal, 8*(1).

Pornpanomchai, C., & Srisupornwattana, N. (2013). Buddhist amulet coin recognition by genetic algorithm. *Proceeding of the 2013 International Computer Science and Engineering Conference* (pp. 324-327). doi:10.1109/ICSEC.2013.6694802.

Pornpanomchai, C., Wongkorsub, J., Pornaudomdaj, T. & Vessawasdi, P. (2010). Buddhist amulet recognition system (BARS). *Proceeding of the 2010 Second International Conference on Computer and Network Technology* (pp. 495-499). doi:10.1109/ICCNT.2010.128.

Prowd, I. (2023, October 10). Types of buddha amulets that bring you more chok so you can live with prosperity. Retrieved from https://thesmartlocal.co.th/buddha-amulets-thailand/

Sauthananusuk, T., Charoenlarpnopparut, C., Kondo, T., Bunnum P., & Hirohiko, K. (2014). Thai amulet recognition using simple feature. *Proceeding of the International Conference of Information and Communication Technology for Embedded Systems.*

Shanthababu, S. (2023, October 10). K-Fold cross validation technique and its essential. Retrieved from https://www.analyticsvidhya.com/blog/2022/02/k-fold-cross-validation-technique-and-its-essentials/

Sucontphunt, T. (2021). Geometric surface image prediction for image recognition enhancement. *Smart Computing and Communication. SmartCom 2020. Lecture Notes in Computer Science 12608*. Cham, Switzerland: Springer. Retrieved from https://doi.org/10.1007/978-3-030-74717-6_27

Weera, K, & Thanasin, B. (2015). Robust texture classification using local correlation features for Thai Buddha amulet recognition. *Applied Mechanics and Materials, 781*, 531-53. Retrieved from https://doi.org/10.4028/www.scientific.net/AMM.781.531

Yuxin F., Quan S., Xinggang W., Tiejun H., Xinlong W., & Yue C. (2023). EVA-02: A visual representation for neon genesis. Retrieved from https://doi.org/10.48550/arXiv.2303.11331