Assignment 2-2 (Group 3)

Anderson Kaina[1] Tanaka Haruki[1] Araki Ryuto[1] Fujii Ryota[1]

Abstract: Human faces often evoke subjective impressions, such as animal-like traits ("cat-like elegance," "dog-like friendliness"). This study investigates computational modeling of such impressions using principal component analysis and support vector machines (SVM). We utilize the PetFace dataset (13 species, over 1M images), preprocess images to 64×64 grayscale, and reduce dimensions to 29 principal components (80% variance). SVM based on a radial basis function kernel is trained for classification, achieving 66.06% accuracy—well above the random baseline of 7.69%. High F1-scores were observed for 'cat,' 'chimp,' and 'hedgehog,' while 'javasparrow' and 'rabbit' performed poorly. Experiments on human faces highlighted cross-domain limitations, as the model still forced animal classifications. Future improvements include addressing class imbalance, enhancing feature extraction, and applying data augmentation. This study demonstrates the potential of SVM for modeling subjective facial impressions.
Keywords— Machine learning, face recognition, animal analogy

1. Introduction

The human face is a complex canvas of expressions and features, often evoking subjective impressions that extend beyond mere identity. This facial impression plays a vital role in our social perception and communication even in digital humans [1]. Thus, quantifying our facial impression and expression will benefit our digital communications in the B5G/6G era. The animal analogy is one of those facial impressions in human faces, which means that a person's face might resemble a particular animal, leading to common descriptors such as "cat-like elegance" or "dog-like friendliness". Historically, this issue has been studied in the psychological research as the terminology of Scocial Cognition [2]. However, their underlying computational mechanisms remain largely unexplored; thus, computational quantification and modeling of these subjective similarities broaden fascinating avenues for developing interdisciplinary research and new digital communication systems.

This assignment investigates animal analogies in human faces using two machine learning techniques, *i.e.*, principal component analysis (PCA) and support vector machines (SVM): We use PetFace dataset [3] and analyze the number of principle components and their corresponding feature maps. Then, we use SVMs based on radial basis function (RBF) kernel for classifying human faces into animal faces, and investigate applicability of SVMs to our facial classification task.

2. Related Work

2.1 PCA and SVM

PCA and SVMs have long been fundamental techniques in the face recognition research due to their robustness and effectiveness [4]. PCA serves as a dimensionality reduction technique, effectively transforming high-dimensional pixel data from facial images into a smaller, more manageable set of principal components that capture the most significant variations in facial features. These components, often referred to as "eigenfaces" in the context of facial analysis, represent the fundamental building blocks of face morphology. Following feature extraction by PCA, SVMs are employed as a robust classification algorithm. SVMs are well-suited for high-dimensional data and are adept at finding an optimal hyperplane that separates different classes with the largest margin, making them ideal for classifying human faces into predefined animal categories based on their extracted features.

The computational analysis of human faces has been a cornerstone of computer vision for decades, leading to significant advancements in face recognition, expression analysis, and attribute prediction. More recently, researchers have begun to delve into the complex task of modeling subjective human impressions of faces, alongside efforts in animal facial analysis. This section reviews relevant literature, laying the groundwork for our proposed approach.

2.2 Machine Learning in Subjective Facial Attributes

The challenge of translating abstract human perceptions into quantifiable computational models has garnered increasing attention. Early works in face perception often relied on psychological studies to understand how humans interpret facial features [4]. With the advent of machine learning, efforts have shifted towards building algorithms that can predict or even manipulate perceived attributes like trustworthiness, attractiveness, or dominance from fa-

[1] Graduate School of Information Science and Technology, The University of Osaka

cial images. For instance, some studies have explored advanced machine learning techniques and generative models to learn latent representations of faces that correlate with these subjective traits [5], [6]. By manipulating these latent spaces, researchers can synthesize faces that embody specific perceived attributes, demonstrating the feasibility of computationally understanding complex human impressions. While these advanced techniques offer powerful capabilities, our study focuses on the more interpretable and computationally efficient combination of PCA and SVM for a specific subjective impression: animal resemblance.

3. Method

3.1 PetFace Datasets

we use PetFace dataset for conducting PCA and implementing SVM, which is a comprehensive resource for animal face identification encompassing 257,484 unique individuals across 13 animal families and 319 breed categories, including both experimental and pet animals (a total of 1,012,934 images). Because the dataset also includes fine-grained annotations such as sex, breed, color, and pattern, the scale and diversity offer significant potential for enhancing future animal-face related research.

In our assignment, input images from the PetFace dataset are first cropped to 64×64 picels around the center, and are converted into gray scale images. Then, they are flattened into a one-dimensional vector of 4096 dimensions, yielding the golobal data matrix whose dimension is $D \times 4096$, where $D = 2000 \times 13$ is the total number of input images as we use 2000 images for each animal class.

3.2 Classification based on PCA and SVM

We applied PCA for dimensionality reduction in order to handle high-dimensional image data efficiently, because it can reduce computational complexity by eliminating redundant dimensions and can prevent overfitting by focusing on the most important variations in the data. The 4096-dimensional image vectors are reduced to 29 principal components which is the minimum number of principal components whose cumulative contribution ratio surpasses 80% of the variance in the dataset. These reduced features are then used for classification with a SVM using a RBF kernel to allow for nonlinear separation between animal classes.

4. Results

4.1 PCA Visualization and Analysis

Principal Component Analysis (PCA) was performed on the flattened image data, producing principal components that capture the variance in the dataset. Figure 1 shows the comparison between original and reconstructed images using 29 principal components, demonstrating that even with fewer dimensions, the reconstructed images retain key features of the original animal faces.

Figure 2 shows the visualization of principal components (eigenvectors). Each component can be interpreted as a pattern that, when combined with others, can reconstruct the

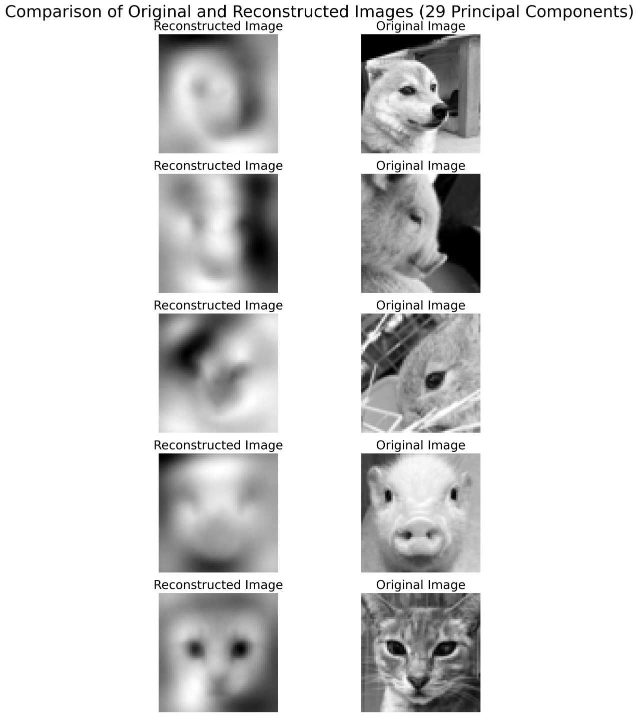

Fig. 1 Comparison of original images (top) and their reconstructions (bottom) using 29 principal components

original faces.

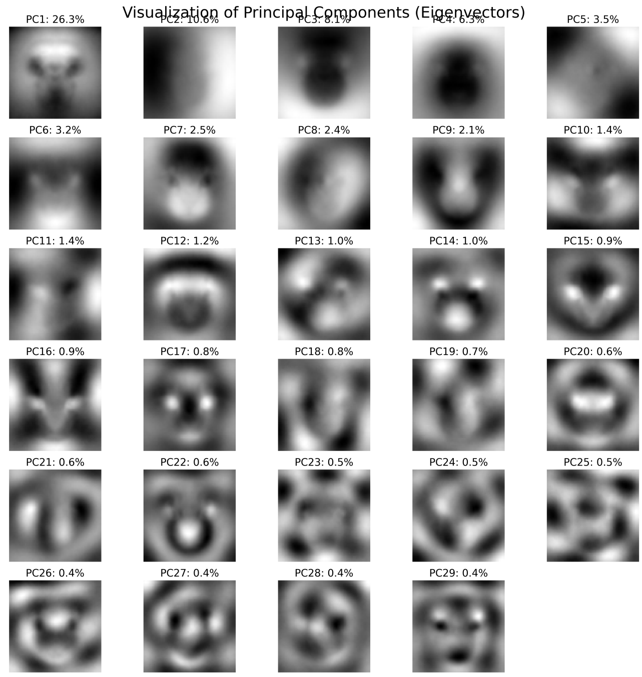

Fig. 2 Feature maps of the first 25 principal components (eigenvectors)

Figure 3 visualizes the data points projected onto the first two principal components, showing the separation of different animal species in the PCA space.

Figure 4 shows the pairwise correlations between the first six principal components, with data points colored by animal species. The diagonal shows the distribution of each principal component by species, while the off-diagonal plots show the relationship between pairs of principal components.

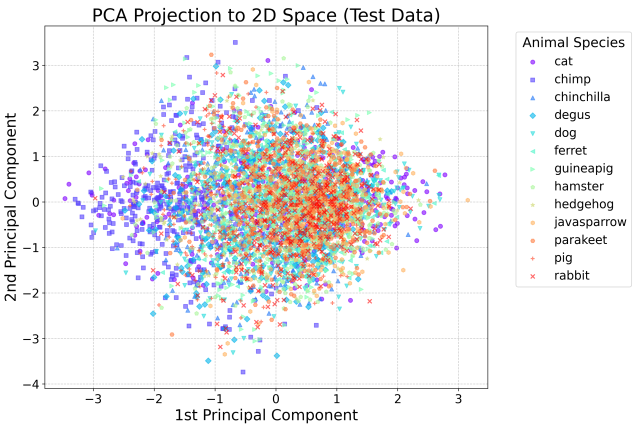

Fig. 3 Projection of test data onto the first two principal components, colored by animal species

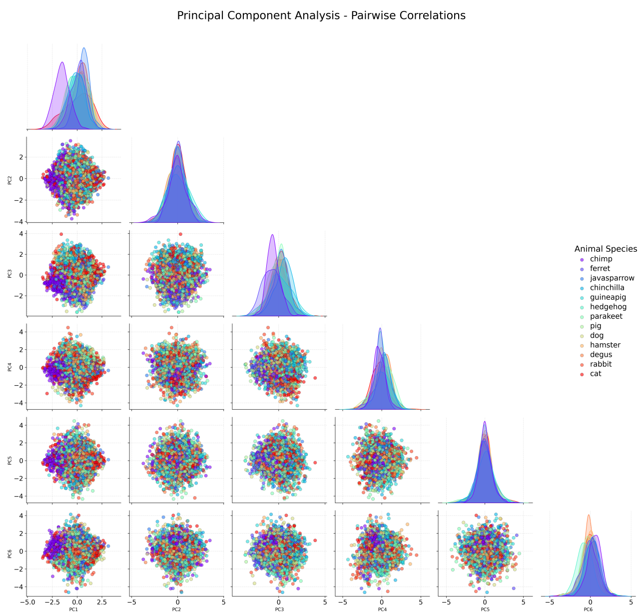

Fig. 4 Pairwise correlations between the first six principal components, colored by animal species

4.2 Precision in SVM

To evaluate the performance of SVM, we split D images into a training set (75%) and a test set (25%). Then, the quantitative evaluation results for the SVM on the test data are presented in Table 1. The model achieved an overall accuracy of 66.06%, which represents a substantial improvement over the theoretical baseline for random classification in a thirteen-class problem (7.69%). This performance indicates moderately effective classification capabilities. Notably, the highest F1-scores were achieved for the 'cat' (0.79), 'chimp' (0.83), and 'hedgehog' (0.79) classes, while the lowest performance was observed for the 'javasparrow' (0.52) and 'rabbit' (0.53) classes, suggesting differential discriminative power across taxonomic categories.

4.3 Classification Results for Target Images

Table 2 shows the results of classifying 6 male face images using the PCA+SVM model trained exclusively on animal face images. The model assigned each human face to an animal class, with chinchilla being the most common classification (3 images), followed by dog, ferret, degus, and hamster (1 image each). This cross-domain application highlights how facial features in humans may share structural similarities with certain animal species, despite the significant domain gap. While these results lack definitive semantic interpretation, they demonstrate the model's ability to extract generalized facial features that translate across species boundaries.

Table 1 Classification Performance Metrics of PCA+SVM

Class	Precision*	Recall**	F1-Score***
cat	0.83	0.76	0.79
chimp	0.78	0.89	0.83
chinchilla	0.57	0.60	0.59
degus	0.57	0.61	0.59
dog	0.73	0.73	0.73
ferret	0.77	0.77	0.77
guineapig	0.62	0.64	0.63
hamster	0.58	0.55	0.56
hedgehog	0.82	0.76	0.79
javasparrow	0.54	0.51	0.52
parakeet	0.58	0.62	0.60
pig	0.64	0.65	0.65
rabbit	0.55	0.51	0.53
Accuracy****			**0.6606**

* The proportion of true positives among all samples predicted as positive.
** The proportion of true positives among all actual positive samples.
*** The harmonic mean of Precision and Recall.
**** The proportion of correctly classified samples.

Table 2 Classification Results for Target Human Face Images

No.	Facial feature	Classified category
1	male_long_beard	ferret
2	male_long	chinchilla
3	male_short_beard	chinchilla
4	male_long_beard	dog
5	male_short_beard	degus
6	male_short	chinchilla
7	male_baby	hamster

Figure 5 presents three examples of the human faces and their corresponding animal classification. Notable patterns emerged: faces with longer beards tended to be classified as dogs or ferrets, while shorter-haired faces were predominantly classified as chinchillas, also round-face were classified as hamster. These associations may be related to the texture and structural patterns that the PCA features extracted from the images.

5. Conclusion

In this assignment, we implemented PCA and SVM for classifying 13 categories of animal images. The experimental results demonstrate 66.06% accuracy for the PCA+SVM approach. This performance substantially exceeds the theoretical random classification baseline of 7.69%, indicating the model's ability to effectively capture discriminative features across multiple animal species. The classification performance varied notably across different taxonomic categories, with certain classes (cat, chimp, hedgehog) achieving F1-scores approaching 0.80, while others (javasparrow, rabbit)

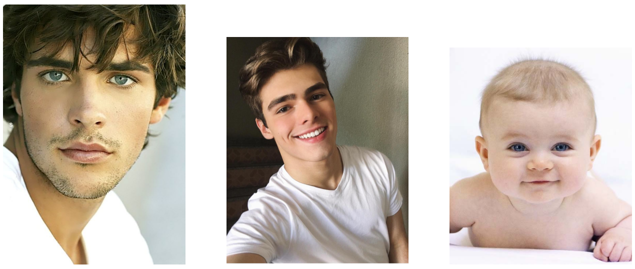

(a) Classified as: dog (b) Classified as: chinchilla (c) Classified as: hamster

Fig. 5 Examples of human face images and their animal classifications by the PCA+SVM model

yielded more modest results around 0.52-0.53. Furthermore, the experiment involving the classification of human face images using a model trained exclusively on animal imagery highlighted the unpredictable behavior of machine learning models when applied to out-of-domain data distributions.

While the achieved accuracy of 66.06% represents a promising result for a complex thirteen-class problem, further improvements could be achieved by addressing several methodological and dataset-specific factors:

(1) **Class Imbalance:** The performance disparity across different animal classes (with F1-scores ranging from 0.52 to 0.83) suggests potential benefits from addressing class representation imbalances or adopting class-specific feature extraction strategies.

(2) **Feature Extraction Refinement:** Though effective, the PCA-based dimensionality reduction approach may not capture all discriminative features necessary for optimal classification performance, particularly for taxonomically similar species.

(3) **Absence of Data Augmentation:** The experimental protocol did not incorporate data augmentation techniques (e.g., rotation, flipping), which could further enhance the training data diversity and improve model generalization capabilities, particularly for the lower-performing classes.

Based on these observations, future research directions could address the following aspects:

- **Implementation of Data Augmentation:** Systematically applying data augmentation strategies to the training dataset to enhance model robustness against variations in pose, illumination, and scale.
- **Advanced Feature Engineering:** Exploring alternative feature extraction methodologies that may more effectively capture the discriminative characteristics of animal facial morphology.
- **Hyperparameter Optimization:** Conducting systematic optimization of hyperparameters such as kernel parameters in SVM and the optimal number of retained principal components to improve classification performance.
- **Domain-Appropriate Evaluation:** To comprehensively assess model performance, utilizing test data from the same distributional domain as the training data, while developing specialized approaches for cross-domain applications.

This study has provided a comparative analysis of basic image classification methodologies and elucidated their inherent limitations. Future research efforts will focus on addressing these identified challenges to develop more accurate and robust image classification systems for both within-domain and cross-domain applications.

Author contribution

R. Araki implemented PCA and SVM. R. Fujii collected datasets and drafted the manuscript. K. Anderson also drafted the manuscript during his international internship. H. Tanaka facilitated group meetings and provided progress reports during the lectures. All authors read and approved the final manuscript.

References

[1] Kikuchi, Shiori, Hisayoshi Ito, and Oky Dicky Ardiansyah Prima. "A Study on the Quality of Facial Expression in Digital Humans."

[2] Zebrowitz, Leslie A., et al. "Animal analogies in first impressions of faces." *Social cognition* 29.4 (2011): 486-496.

[3] Shinoda, Risa, and Kaede Shiohara. "Petface: A large-scale dataset and benchmark for animal identification." *European Conference on Computer Vision*. Cham: Springer Nature Switzerland, 2024.

[4] M. Turk and A. Pentland, "Eigenfaces for Recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.

[5] P. Salehi, S. A. Sheshkal, V. Thambawita, and P. Halvorsen, "FROM FLAT TO FEELING: A FEASIBILITY AND IMPACT STUDY ON DYNAMIC FACIAL EMOTIONS IN AI-GENERATED AVATARS," arXiv preprint arXiv:2506.13477, 2025.

[6] J. C. Peterson, S. Uddenberg, T. L. Griffiths, A. Todorov, and J. W. Suchow, "Deep models of superficial face judgments," *Proceedings of the National Academy of Sciences (PNAS)*, vol. 119, no. 17, p. e2115228119, 2022.