

COMP64301: Cognitive Robotics and Computer Vision Assignment

Zhikun Peng (14247379)

December 3, 2025

1 Introduction

1.1 Background

Object recognition is a basic ability for cognitive robots that enables them to communicate with their surrounding environments via navigation, manipulation, and human-robot interaction (HRI) [1]. Previously, this area of research focused on the use of hand-crafted feature extraction techniques like the Scale-Invariant Feature Transform (SIFT) [2] and Histogram of Oriented Gradients (HOG), commonly aggregated in the Bag of Visual Words (BoW) model [3]. Although these methods provided interpretability and robustness to geometric transformations, they often struggle with high-level semantic variations.

In recent years, the new generation of deep learning methods – mainly focusing on CNNs – have become an emerging trend that outperforms many previous models. Models such as ResNet [4] have revolutionized computer vision by learning hierarchical feature representations directly from data, achieving state-of-the-art performance on large-scale benchmarks like ImageNet. But understanding how the classical and novel solutions match is still necessary to deploy efficient vision systems on the limited resources robotics platforms.

1.2 Research Objectives

The main goal of this report is to analyse the effectiveness of Deep Learning (ResNet18) and Traditional Computer Vision (SIFT+BoW). The research will concentrate on evaluating how well feature learning (CNN) and feature engineering (BoW) perform at different levels of difficulty. The impact of hyperparameters such as vocabulary size (BoW), batch size (CNN), etc., on classification accuracy and computation efficiency were explored systematically. Then the transfer learning's generalization capability on domain-specific small dataset was analyzed.

1.3 Introduction to the Dataset

To strictly evaluate the performance robustness of the mentioned algorithm, two kinds of datasets are utilized to simulate various appearance problems existing in robotic vision:

A subset of the Caltech 101 dataset [5] is used as the experiment benchmark. It includes various categories with relatively large inter-class differences, such as Airplanes, Motorbikes, Bonsai, and so on. The images belonging to different classes have different structures, making it a relatively coarse-grained classification. This dataset can be taken to verify whether the algorithm is correct.

The sub-dataset of the Oxford-IIIT Pet data set [6] was selected for the more challenging task (representing a finer classifying task), which included dogs and cats of different breeds, characterized by large intra-class variances (e. g., pose, scale) and small inter-class variances (e. g., similar fur texture and shape). As such, this was more analogous to a complex real-world scenario where a robot had to identify visually similar objects.

2 Methodology

Study used an effective comparative scheme and examined both feature extraction methods: traditional feature engineering (BoW) and modern feature learning using CNNs with transfer learning. It also made sure that all the images were

divided up before they were saved as train, validation and test splits in a ratio of 70:15:15. A fixed random seed of 42 was applied during this operation.

2.1 Deep Learning Approach: Convolutional Neural Networks (CNN)

2.1.1 Model Selection and Architecture

The ResNet18 architecture [4] was chosen for its efficient convergence and robust performance. It offers a favorable balance between computational efficiency on the NVIDIA RTX 4060 GPU and high performance for image classification tasks.

2.1.2 Transfer Learning Strategy

Since the dataset remains too small to generate ideal results, a transfer learning strategy is adopted to prevent overfitting and enhance computational efficiency. For instance, ImageNet pre-trained weights were used for initialization, with all convolutional backbone layers frozen to function as a fixed and reliable feature extractor. Only the final fully connected layer required optimization to map extracted features to ten target categories. The parameters of this fully connected layer were reset, enabling it to output the desired ten-class predictions.

2.1.3 Data Preprocessing and Optimisation

All input images were resized to 224×224 pixels, converted into PyTorch tensors, and normalized using the standard ImageNet values. To enhance generalization, random horizontal flipping was applied as a data augmentation technique during training. The training process employed the Adam optimizer in conjunction with the Cross-Entropy loss function. Additionally, systematic hyperparameter exploration was conducted to ensure the efficiency and stability of the convolutional neural network (CNN). Specifically, convergence behavior was evaluated by testing learning rates $\{0.001, 0.0005\}$, while batch sizes $\{32, 64\}$ were tested to analyze the trade-off between gradient stability and GPU utilization.

2.2 Traditional Approach: Bag of Visual Words (BoW)

2.2.1 Feature Extraction and Encoding

For Caltech101 and Oxford Pets datasets where objects' scales, orientations and illuminations vary greatly, SIFT [2] is chosen as local feature descriptor because of its robustness in facing aforementioned problems. All SIFT descriptors in the training set are pooled to learn a visual dictionary for describing images later on, which is represented by histograms after encoding every image.

2.2.2 Visual Vocabulary Construction and Hyperparameter Tuning

The visual vocabulary was constructed by clustering the aggregated SIFT descriptors using the K-Means algorithm. The number of clusters, denoted as K , represents the vocabulary size and serves as a critical hyperparameter governing the trade-off between discriminative power and computational efficiency. To quantify its impact, we systematically investigated $K \in \{50, 100, 200, 500\}$.

2.2.3 Classification via Support Vector Machines (SVM)


Using an SVM classifier based on radial basis functions (RBF), with BoF histograms as input. Two kernel functions (linear kernel and radial basis function kernel) were compared, with the RBF kernel ultimately selected as the subject of study for its ability to more accurately describe the distribution of nonlinear features. The regularization coefficient C was tested within the range $\{0.1, 1.0, 10.0\}$ to balance maximizing classification error with penalizing misclassified training points.

3 Experimental Setup and Results


3.1 Experimental Configuration

The dataset was stratified by category proportion for training, validation, and testing sets to maintain balanced class distribution across each set. The final distribution was: training set 70%, validation set 15%, and testing set 15%. Model evaluation employed two primary metrics: test accuracy (%) as an effectiveness indicator and training time (seconds) as an efficiency indicator. All convolutional neural networks (CNNs) were trained on NVIDIA RTX 4060 GPUs, while the Bow experiment utilized only CPUs. From the Caltech101 dataset, the following 10 objects were selected:airplanes, motorbikes, faces_easy, watch, leopards, bonsai, car_side, ketch, chandelier, and hawkbill. In addition, ten cat or dog breeds were drawn from the Oxford Pet dataset, namely Abyssinian, Bengal, Birman, Bombay, British Shorthair, Egyptian Mau, Maine Coon, Persian, Ragdoll, and Russian Blue.


3.2 Results: CNN Hyperparameter Exploration




(a) Caltech101: Confusion Matrix



(b) Oxford Pets: Confusion Matrix



(c) Oxford Pets: Loss & Accuracy



(d) Oxford Pets: Loss & Accuracy

Figure 1: Visualization of CNN performance with Learning Rate=0.001, Batch Size=64 and Epochs=30. The (a) and (b) show the Confusion Matrices, where the diagonal represents correct predictions. The (c) and (d) show the Training/Validation Loss and Accuracy over epochs.

The confusion matrix in Figure 1 (a) exhibits a perfect diagonal distribution with no off-diagonal elements. This indicates the model perfectly distinguishes coarse-grained categories like airplanes and motorcycles, demonstrating that ImageNet pre-trained features exhibit “overfitting” characteristics for such simple tasks. According to Figure 1 (c), ResNet18 converges extremely rapidly on the Caltech101 dataset. By the third epoch, validation accuracy had already saturated near 100%, and the loss function rapidly dropped to zero.

The situation for Oxford Pets is slightly different. Although the final accuracy is high (around 90%+), as shown in Figure 1 (d), the loss decreases more gradually than Caltech, and a slight gap exists between the training and validation curves, suggesting a minor risk of overfitting. The confusion matrix in Figure 1 (b) reveals specific errors: the model occasionally confuses visually similar breeds (e.g., “Russian Blue” and “British Shorthair”).

The performance of the ResNet18 model under parts of various hyperparameter configurations is summarized in Table 1. The results indicate a significant dichotomy in task difficulty between the two datasets.

Table 1: Selected CNN Experimental Results. The table contrasts the saturation on the coarse-grained task (Caltech101) with the hyperparameter sensitivity on the fine-grained task (Oxford Pets).

Dataset	Batch Size	LR	Epochs	Test Acc (%)	Time (s)
<i>Caltech101</i>	64	0.001	15	99.78	76
<i>Caltech101</i>	32	0.001	30	100.00	166
<i>Oxford Pets</i>	64	0.0005	15	88.33	116
<i>Oxford Pets</i>	32	0.001	30	90.67	234
<i>Oxford Pets</i>	32	0.0005	30	92.67	225

On the Caltech101 dataset in Table 1, the model performed nearly flawlessly across all settings, achieving a maximum accuracy of 100%. Even when training rounds were reduced from 30 to 15, accuracy remained remarkably high at 99.78%. This demonstrates that the pre-trained ResNet excels at extracting high-level features (such as the shapes of airplanes or motorcycles), with hyperparameters having minimal impact on results. The model is sufficiently robust to accomplish this task.

Even though there were only 30 images in each class of the training data, the model did not suffer too much from overfitting. The "frozen backbone" method is attributed for this kind of stability because it constrained the tunable parameters in our CNN model to the last fully-connected layer, a useful regularizer for this work.

The Oxford Pet Dataset presents greater challenges due to high inter-class similarity, with models exhibiting high sensitivity to hyperparameters (see Table 1). Using a lower learning rate of 0.0005 paired with a batch size of 32 achieved the highest accuracy of 92.67%, surpassing the 90.67% accuracy obtained with a learning rate of 0.001. This indicates that fine-grained classification tasks benefit from smaller step updates, which optimize decision boundaries and reduce overfitting. Additionally, the dataset requires 30 training epochs to fully converge, representing a roughly 4% improvement over 15 epochs. In terms of efficiency, GPU acceleration delivers outstanding performance: Caltech101 converged in just 76 seconds, while the Oxford Pets experiment completed in under four minutes.

3.3 Results: BoW Parameter Evaluation

3.3.1 Parameter Sensitivity on Caltech101




Figure 2: BoW Performance on Caltech101: Test Accuracy (%) vs. Vocabulary Size (K) for Linear and RBF SVM Kernels.


Figure 2 describes the Caltech 101 performance against two of the key hyperparameters, Vocabulary Size (K) and SVM Kernel. The experiments demonstrate that there is a high positive correlation between Vocabulary Size (K) and classification accuracy; the increase in K caused the increase in classification accuracy of 77.46% when K was 50(best accuracy) and 85.34% when K increased to 500(best accuracy).

The Radial Basis Function (RBF) kernel produced better results compared to Linear kernel with on average 3%-5% of the difference. In view of this, Bag of Words is represented as a higher dimensional feature space, thus suited to capturing


the inherent non-linear and complex relationship of High dimensional BoW space model, and therefore having more advantages than Linear kernel which cannot handle more complex functions.

When regularization coefficient (C) is set to very small value such as $C = 0.1$, the resulting prediction accuracy is quite low; such as only 56% accuracy for some K choices. This fact reflects an under-fitting situation when the penalty on model complexity is over-strong. On the contrary, the best model can be found when combining $C=10.0$ and RBF kernel to balance between maximizing the margin and allowing a few misclassified training points.

3.3.2 Analysis of Dataset Difficulty



(a) Caltech101: Confusion Matrix



(b) Oxford Pets: Confusion Matrix

Figure 3: Visualization of BoW performance with $K=100$, $C=1.0$ and Kernel=RBF. The (a) and (b) show the Confusion Matrices, where the diagonal represents correct predictions.

The confusion matrix for Oxford Pets (Figure 3 (b)) reveals that even with the optimal parameter combination ($K = 500$, RBF kernel), the bag-of-words model achieves a maximum accuracy of only 41.67% on the Oxford Pets dataset. The off-diagonal region exhibits numerous scattered misclassification points. This indicates the model struggles significantly to distinguish between different breeds of cats and dogs. The confusion is not confined to specific breed pairs but reflects a general inability to extract discriminative features.

4 Comparison and Discussion

4.1 Performance Comparison: Feature Engineering vs. Feature Learning

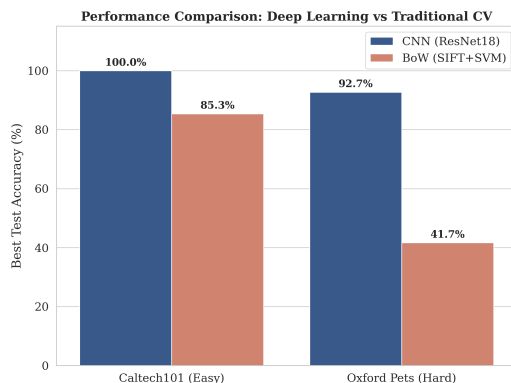


Figure 4: Performance Comparison between Deep Learning (CNN) and Traditional CV (BoW) across datasets. While both methods perform adequately on the coarse-grained task, a massive performance gap emerges on the fine-grained task.

A comparison of top scoring approaches as shown in Figure 4, the significant differences brought by the deep learning paradigm are clearly evident. Overall, on the coarse-grained Caltech101 dataset, performance variations exist but are not

substantial. While CNN achieves 100% accuracy, the BoW method still attains 85.3%. This suggests that for discriminative objects like those depicted here, manually designed features (such as SIFT) are sufficient to learn object patterns for retrieval.

While there is a difference in the coarse structure, the difference is huge when we consider the fine-grained Oxford Pet dataset. While the CNN still retains the accuracy at 92.7%, while the BoW's performance plummets to 41.7%. This gap is because the BoW model only pays attention to the frequency of occurrence of local features (e. g., corners, edges) and does not account for the location of these keypoint features, nor do they capture the semantics of each other. When classes are too similar (texture and shape wise), like fur, the "bag of words" cannot tell them apart. Instead, the hierarchical layers of ResNet18 can effectively abstract the high-level semantic features like ear shape, snout structure that serve as cues for breed discrimination.

4.2 Computational Efficiency and Resource Utilization

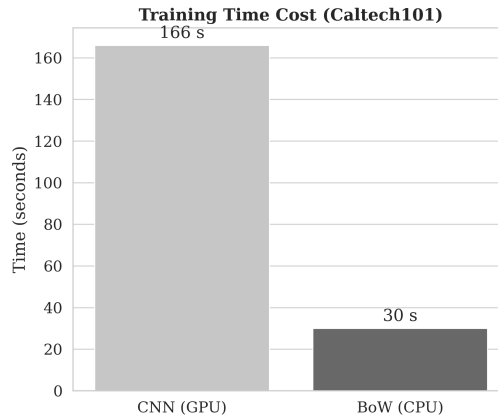


Figure 5: Training Time Cost Comparison on Caltech101. CNN training utilized an NVIDIA RTX 4060 GPU, while BoW training was performed on a AMD Ryzen 7 8845H CPU.

Figure 5 describes the trade-off between accuracy and computational costs, we can observe that the BoW model trains relatively faster for the Caltech101 dataset, it takes around 30 seconds to be trained on the CPU only. Therefore, for very constrained scenarios where GPUs are not available (like some low-power microcontrollers), this model is still computationally appealing because of its low-computational requirements.

Contrariwise, the CNN needed 166 sec. (for the 30-epoch run) and was run using an NVIDIA RTX 4060 GPU. Though the actual training time for CNN is much more, the drastic reduction through the help of transfer learning should be noted here. It would have taken hours, even days, to train a brand new deep network from scratch, while simply finetuning the output layer made our CNN achieve almost state-of-the-art accuracy within three minutes of training under the modern cognitive robotic system with widespread applications of GPUs (e.g., NVIDIA Jetson).

4.3 Summary of Trade-offs

In summary, CNNs achieve high accuracy and robustness through automatic feature learning, but they rely on hardware acceleration and offer limited interpretability. In contrast, BoW methods are lightweight, intuitive, and interpretable, yet their performance is constrained in fine-grained classification due to the absence of spatial encoding, and they are sensitive to parameter selection. This creates a trade-off: CNNs sacrifice computational resources for superior performance, while BoW methods are efficient but struggle with complex visual tasks.

5 State of the Art: Contextualisation in Robotics

Although convolutional neural networks excel in controlled classification tasks (Section 4), their application to autonomous systems requires addressing broader perception challenges beyond simple recognition. This section contextualizes our findings within the forefront of robot vision research.

Historically, robotics relied on geometric features (SIFT/ORB) for tasks like SLAM [7]. While robust for mapping, these methods lacked the semantic understanding required for interaction. Deep learning revolutionized this landscape by enabling semantic perception. As demonstrated experimentally, convolutional neural networks extract high-level abstract information, empowering robots to perform semantic navigation and grasping tasks [8].

However, the field is rapidly shifting toward embodied AI and multimodal learning. Unlike passive recognition, embodied AI posits that intelligence arises from physical interactions between agents and their environment [9]. Recent advancements like CLIP and RT-2 (Robot Transformer 2) [10] fuse visual data with natural language instructions through visual-language models (VLMs). This multimodal approach enables robots to generalize to unseen objects (zero-shot learning) and execute complex commands like “pick up the extinct animal toy” (referring to a dinosaur)—a breakthrough unattainable by monomodal convolutional neural networks or bag-of-words models.

Despite these advances, deploying large foundational models on mobile platforms remains constrained by stringent SWaP (size, weight, and power) limitations. Consequently, current research focuses on transfer learning from simulated to real environments and efficient model distillation techniques to bridge the gap between high-level semantic reasoning and real-time physical control.

6 Conclusion

This report compares Deep Learning and traditional Computer Vision approaches in object recognition of cognitive robots, where we empirically demonstrated that SIFT+BoW is adequate for coarse-grained tasks; but it cannot be adapted to finer-grained recognition tasks. By comparison, using transfer learning, the ResNet18-based CNN achieves better results steadily compared with other models, which indicates that learning features outperforms the hard-coded ones. In sum, Deep Learning is deemed to be an advantageous choice when handling a much greater variety of situations encountered by autonomous agents in challenging environments, but a future work focus will probably be more emphasis on seeking approaches to lighten up the learning process, such as the use of MobileNet architecture to tackle smaller errors and trade-off between higher computational cost.

References

- [1] A. Cangelosi and M. Asada, “What is cognitive robotics?” In *Cognitive Robotics*, The MIT Press, May 2022, ISBN: 9780262369329. DOI: 10.7551/mitpress/13780.003.0005. eprint: https://direct.mit.edu/book/chapter-pdf/2239481/c000600_9780262369329.pdf. [Online]. Available: <https://doi.org/10.7551/mitpress/13780.003.0005>.
- [2] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [3] G. Csurka, C. R. Dance, L. Fan, J. Willamowski, and C. Bray, “Visual categorization with bags of keypoints,” *workshop on statistical learning in computer vision eccv*, 2004.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *IEEE*, 2016.
- [5] F.-F. Li, M. Andreeto, M. Ranzato, and P. Perona, *Caltech 101*, Apr. 2022. DOI: 10.22002/D1.20086.
- [6] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar, “Cats and dogs,” in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3498–3505. DOI: 10.1109/CVPR.2012.6248092.
- [7] C. Cadena et al., “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016. DOI: 10.1109/TR0.2016.2624754.
- [8] N. Sünderhauf et al., “The limits and potentials of deep learning for robotics,” *The International journal of robotics research*, vol. 37, no. 4-5, pp. 405–420, 2018.
- [9] J. Duan, S. Yu, H. L. Tan, H. Zhu, and C. Tan, “A survey of embodied ai: From simulators to research tasks,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 6, no. 2, pp. 230–244, 2022.
- [10] B. Zitkovich et al., “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” in *Conference on Robot Learning*, PMLR, 2023, pp. 2165–2183.

A Appendix: Source Code

A.1 Data Splitting Script

```
1 #generate_splits_and_folders.py
2 import os
3 import glob
4 import pandas as pd
5 import numpy as np
6 import shutil
7 from sklearn.model_selection import train_test_split
8 from tqdm import tqdm
9
10 # ===== Configuration Area =====
11
12 # Original data root directory
13 RAW_DATA_DIR = r"D:\comp64301_cv\data\raw"
14 # Output directory (processed data will be saved here)
15 OUTPUT_DIR = r"D:\comp64301_cv\data\processed"
16
17 # 1. Caltech101 configuration
18 CALTECH_DIR = os.path.join(RAW_DATA_DIR, "caltech-101")
19 CALTECH_TARGET_CLASSES = [
20     "airplanes", "motorbikes", "Faces_easy", "watch", "leopards",
21     "bonsai", "car_side", "ketch", "chandelier", "hawksbill"
22 ]
23
24 # 2. Oxford-IIIT Pet configuration
25 OXFORD_DIR = os.path.join(RAW_DATA_DIR, "oxfordiiit-pet", "images")
26 OXFORD_CLASS_COUNT = 10
27
28 # Split ratios (must sum to 1.0)
29 TRAIN_RATIO = 0.70
30 VAL_RATIO = 0.15
31 TEST_RATIO = 0.15
32 # =====
33
34 def ensure_dir(directory):
35     if not os.path.exists(directory):
36         os.makedirs(directory)
37
38 def get_caltech_data():
39     """Get Caltech101 data list"""
40     data = []
41     print(f"Scanning Caltech101: {CALTECH_DIR}...")
42
43     # Note: Handle case sensitivity issues here, try to match folders
44     if not os.path.exists(CALTECH_DIR):
45         print(f"Warning: Path not found {CALTECH_DIR}")
46         return pd.DataFrame()
47
48     # Get the list of actual existing folders
49     existing_folders = os.listdir(CALTECH_DIR)
50     # Create a case-insensitive mapping dictionary
51     folder_map = {f.lower(): f for f in existing_folders}
52
53     for target_cls in CALTECH_TARGET_CLASSES:
54         # Try to match
55         actual_folder = folder_map.get(target_cls.lower())
56
57         if actual_folder:
58             class_path = os.path.join(CALTECH_DIR, actual_folder)
59             images = glob.glob(os.path.join(class_path, "*.jpg"))
60             for img_path in images:
61                 data.append({
62                     "filepath": img_path,
63                     "label": target_cls, # use original target class name
64                     "dataset_source": "caltech101"
65                 })
66         else:
67             print(f"Warning: Folder for class '{target_cls}' not found in Caltech101")
68
69     print(f"Caltech101: Found {len(data)} images.")
```

```

70     return pd.DataFrame(data)
71
72 def get_oxford_data():
73     """Get Oxford Pets data list"""
74     data = []
75     print(f"Scanning Oxford Pets: {OXFORD_DIR}...")
76
77     if not os.path.exists(OXFORD_DIR):
78         print(f"Warning: Path not found {OXFORD_DIR}")
79         return pd.DataFrame()
80
81     all_images = glob.glob(os.path.join(OXFORD_DIR, "*.jpg"))
82
83     temp_list = []
84     for img_path in all_images:
85         filename = os.path.basename(img_path)
86         # Oxford dataset filenames are usually Class_Name_Number.jpg
87         # Method: find the last underscore and take the part to the left of it
88         class_name = "_".join(filename.split("_")[:-1])
89         temp_list.append((img_path, class_name))
90
91     df = pd.DataFrame(temp_list, columns=["filepath", "label"])
92
93     # Select top N classes in alphabetical order
94     unique_classes = sorted(df["label"].unique())
95     target_classes = unique_classes[:OXFORD_CLASS_COUNT]
96     print(f"Oxford Pets selected top {OXFORD_CLASS_COUNT} classes: {target_classes}")
97
98     # Filter data
99     df_filtered = df[df["label"].isin(target_classes)].copy()
100    df_filtered["dataset_source"] = "oxford_pet"
101
102    print(f"Oxford Pets: Found {len(df_filtered)} images.")
103    return df_filtered
104
105 def _safe_train_test_split(df, test_size, stratify_col="label", random_state=42):
106    """Try stratified sampling, fallback to random sampling if fails (e.g., too few samples in a
107    class)"""
108    stratify = df[stratify_col] if stratify_col in df.columns else None
109    try:
110        return train_test_split(
111            df, test_size=test_size, random_state=random_state, stratify=stratify
112        )
113    except ValueError as e:
114        print(f"Warning: Stratified sampling failed, falling back to random sampling. Reason: {e}")
115    return train_test_split(
116        df, test_size=test_size, random_state=random_state, stratify=None, shuffle=True
117    )
118
119 def organize_files_into_folders(df, dataset_name):
120    """
121    Based on the split column in the DataFrame, physically copy images into folder structures.
122    Target structure: output_dir / dataset_name / class_name / image.jpg
123    """
124    print(f"\nOrganizing file structure and copying images to: {os.path.join(OUTPUT_DIR,
125    dataset_name)} ...")
126
127    # Iterate over each row in the DataFrame
128    # Use tqdm to show progress bar
129    for _, row in tqdm(df.iterrows(), total=len(df), desc=f"Processing {dataset_name}"):
130        src_path = row['filepath']
131        split = row['split'] # train, val, test
132        label = row['label'] # class name
133
134        # Construct target directory
135        # Example: D:\comp64301_cv\data\processed\caltech101\train\airplanes\
136        target_dir = os.path.join(OUTPUT_DIR, dataset_name, split, label)
137
138        # Ensure target directory exists
139        os.makedirs(target_dir, exist_ok=True)
140
141        # Construct target filename (to avoid name conflicts, keep original filename)
142        filename = os.path.basename(src_path)

```

```

141     target_path = os.path.join(target_dir, filename)
142
143     # Copy file (copy2 preserves metadata like timestamps)
144     try:
145         shutil.copy2(src_path, target_path)
146     except Exception as e:
147         print(f"Copy failed: {src_path} -> {target_path}, Error: {e}")
148 def process_dataset(df, dataset_name):
149     """Process a single dataset: split, save CSV, and organize folders"""
150     if len(df) == 0:
151         print(f"Error: {dataset_name} dataset is empty, skipping.")
152         return
153
154     print(f"\nProcessing dataset: {dataset_name}")
155     print(f"Using ratios: train={TRAIN_RATIO}, val={VAL_RATIO}, test={TEST_RATIO}")
156     # ===== Splitting logic =====
157     parts = []
158
159     # 1. Split Test set
160     if TEST_RATIO > 0:
161         train_val, test = _safe_train_test_split(df, test_size=TEST_RATIO, stratify_col="label")
162         test = test.copy()
163         test.loc[:, "split"] = "test"
164         parts.append(test)
165     else:
166         train_val = df
167
168     # 2. Split Train / Val
169     remain_ratio = 1.0 - TEST_RATIO
170     if remain_ratio > 0:
171         # Calculate relative proportion of val in the remaining data
172         # relative_val = target_val / (target_train + target_val)
173         if (TRAIN_RATIO + VAL_RATIO) > 0:
174             relative_val_size = VAL_RATIO / (TRAIN_RATIO + VAL_RATIO)
175
176             train, val = _safe_train_test_split(train_val, test_size=relative_val_size,
177 stratify_col="label")
178             train = train.copy()
179             val = val.copy()
180             train.loc[:, "split"] = "train"
181             val.loc[:, "split"] = "val"
182             parts.extend([train, val])
183         else:
184             # Edge case handling
185             train_val.loc[:, "split"] = "train" # Default to train
186             parts.append(train_val)
187
188     # Combine results
189     final_df = pd.concat(parts, ignore_index=True)
190
191     # ===== Save CSV =====
192     csv_filename = f"{dataset_name}_split.csv"
193     csv_path = os.path.join(OUTPUT_DIR, csv_filename)
194     final_df.to_csv(csv_path, index=False)
195     print(f"CSV saved: {csv_path}")
196     print(f"Distribution stats:\n{final_df['split'].value_counts()}")
197
198     # ===== Physical file copying =====
199     organize_files_into_folders(final_df, dataset_name)
200     print(f"{dataset_name} processing completed.")
201
202 if __name__ == "__main__":
203     ensure_dir(OUTPUT_DIR)
204
205     # 1. Caltech101
206     df_caltech = get_caltech_data()
207     process_dataset(df_caltech, "caltech101")
208
209     # 2. Oxford Pets
210     df_oxford = get_oxford_data()
211     process_dataset(df_oxford, "oxford_pet")
212
213     print("\nAll datasets processed successfully.")

```

Listing 1: generate_splits_and_folders.py

A.2 CNN Training Script

```
1 #train_cnn.py
2 import torch
3 import torch.nn as nn
4 import torch.optim as optim
5 from torchvision import datasets, models, transforms
6 from torch.utils.data import DataLoader
7 import os
8 import time
9 import copy
10 import csv
11 from datetime import datetime
12 import matplotlib.pyplot as plt
13 from sklearn.metrics import confusion_matrix
14 import seaborn as sns
15 import numpy as np
16
17 # ===== Setting Area =====
18 # Dataset directories
19 DATA_DIR_CALTECH101 = r"D:\comp64301_cv\data\processed\caltech101"
20 DATA_DIR_oxford_PET = r"D:\comp64301_cv\data\processed\oxford_pet"
21
22 # Model save directory
23 MODEL_SAVE_DIR = r"D:\comp64301_cv\results\cnn_models"
24 # Figure save directory
25 FIGURE_SAVE_DIR = r"D:\comp64301_cv\results\figures"
26 # Experiment results log file
27 LOG_CSV_FILE = r"D:\comp64301_cv\results\cnn_results.csv"
28
29 # Ensure save directories exist
30 os.makedirs(MODEL_SAVE_DIR, exist_ok=True)
31 os.makedirs(FIGURE_SAVE_DIR, exist_ok=True)
32 os.makedirs(os.path.dirname(LOG_CSV_FILE), exist_ok=True)
33
34 # Device (4060 cuda available)
35 DEVICE = torch.device("cuda" if torch.cuda.is_available() else "cpu")
36 # =====
37
38 def get_dataloaders(data_dir, batch_size=32):
39     """
40     Load data and apply normalization.
41     ImageNet normalization parameters: mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]
42     """
43     data_transforms = {
44         'train': transforms.Compose([
45             transforms.Resize((224, 224)), # ResNet input standard size
46             transforms.RandomHorizontalFlip(), # Data augmentation
47             transforms.ToTensor(),
48             transforms.Normalize([0.485, 0.456, 0.406], [0.229, 0.224, 0.225])
49         ]),
50         'val': transforms.Compose([
51             transforms.Resize((224, 224)),
52             transforms.ToTensor(),
53             transforms.Normalize([0.485, 0.456, 0.406], [0.229, 0.224, 0.225])
54         ]),
55         'test': transforms.Compose([ # Test processing usually same as Val
56             transforms.Resize((224, 224)),
57             transforms.ToTensor(),
58             transforms.Normalize([0.485, 0.456, 0.406], [0.229, 0.224, 0.225])
59         ]),
60     }
61
62     image_datasets = {x: datasets.ImageFolder(os.path.join(data_dir, x), data_transforms[x])
63                     for x in ['train', 'val', 'test']}
64
65     # Pointing to 'test' subfolder
66     test_dataset = datasets.ImageFolder(os.path.join(data_dir, 'test'), data_transforms['test'])
67
68     # shuffle=False because no need to shuffle
69     test_loader = DataLoader(test_dataset, batch_size=batch_size, shuffle=False, num_workers=0)
70
71     # On Windows, num_workers=0 is usually the most stable to avoid multi-process errors
72     dataloaders = {x: DataLoader(image_datasets[x], batch_size=batch_size, shuffle=(x!='train'),
```

```

num_workers=0)
    for x in ['train', 'val', 'test']}
73
74
75 dataset_sizes = {x: len(image_datasets[x]) for x in ['train', 'val', 'test']}
76 class_names = image_datasets['train'].classes
77
78 return dataloaders, dataset_sizes, class_names
79
80 def initialize_model(num_classes, feature_extract=True):
81     """
82     Initialize ResNet18 model
83     :param feature_extract: True = freeze backbone, only train FC layer; False = fine-tune entire
84     network
85     """
86     print("Initializing ResNet18...")
87     # Download/load pretrained weights
88     model = models.resnet18(weights=models.ResNet18_Weights.DEFAULT)
89
90     # If feature extracting, freeze parameters
91     if feature_extract:
92         for param in model.parameters():
93             param.requires_grad = False # Key step: disable gradient computation
94
95     # Replace the final fully connected layer (fc)
96     # Newly created layers have requires_grad=True by default
97     num_ftrs = model.fc.in_features
98     model.fc = nn.Linear(num_ftrs, num_classes)
99
100     return model
101
102 def train_model(model, dataloaders, dataset_sizes, criterion, optimizer, num_epochs=10):
103     """
104     Training loop function
105     """
106     since = time.time()
107
108     best_model_wts = copy.deepcopy(model.state_dict())
109     best_acc = 0.0
110
111     # Used to record history (used for plotting)
112     history = {'train_loss': [], 'train_acc': [], 'val_loss': [], 'val_acc': []}
113
114     print(f"Starting training... Device: {DEVICE}")
115
116     for epoch in range(num_epochs):
117         print(f'Epoch {epoch+1}/{num_epochs}')
118         print('-' * 10)
119
120         # Each epoch has a training and validation phase
121         for phase in ['train', 'val']:
122             if phase == 'train':
123                 model.train() # Training mode
124             else:
125                 model.eval() # Evaluation mode
126
127             running_loss = 0.0
128             running_corrects = 0
129
130             # Iterate over data
131             for inputs, labels in dataloaders[phase]:
132                 inputs = inputs.to(DEVICE)
133                 labels = labels.to(DEVICE)
134
135                 # Zero the gradients
136                 optimizer.zero_grad()
137
138                 # Forward propagation
139                 # Only track gradients in the training phase
140                 with torch.set_grad_enabled(phase == 'train'):
141                     outputs = model(inputs)
142                     _, preds = torch.max(outputs, 1)
143                     loss = criterion(outputs, labels)
144
145                 # Backpropagation + optimization (only in training phase)

```

```

145         if phase == 'train':
146             loss.backward()
147             optimizer.step()
148
149         # Statistics
150         running_loss += loss.item() * inputs.size(0)
151         running_corrects += torch.sum(preds == labels.data)
152
153     epoch_loss = running_loss / dataset_sizes[phase]
154     epoch_acc = running_corrects.double() / dataset_sizes[phase]
155
156     print(f'{phase} Loss: {epoch_loss:.4f} Acc: {epoch_acc:.4f}')
157
158     # Record history
159     history[f'{phase}_loss'].append(epoch_loss)
160     history[f'{phase}_acc'].append(epoch_acc.item())
161
162     # If in validation phase and accuracy improves, save model weights
163     if phase == 'val' and epoch_acc > best_acc:
164         best_acc = epoch_acc
165         best_model_wts = copy.deepcopy(model.state_dict())
166
167     time_elapsed = time.time() - since
168     print(f'Training complete in {time_elapsed // 60:.0f}m {time_elapsed % 60:.0f}s')
169     print(f'Best val Acc: {best_acc:.4f}')
170
171     # Load best model weights
172     model.load_state_dict(best_model_wts)
173     # Modify return values: add time_elapsed
174     return best_acc, model, history, time_elapsed
175
176 def evaluate_on_test(model, test_loader, device):
177     """
178     Evaluate model performance on the Test set
179     """
180     model.eval() # Switch to evaluation mode
181
182     correct = 0
183     total = 0
184
185     # To store all predictions and labels
186     all_preds = []
187     all_labels = []
188
189     print("\nPerforming final evaluation on the Test set...")
190
191     # Disable gradient calculation
192     with torch.no_grad():
193         for inputs, labels in test_loader:
194             inputs = inputs.to(device)
195             labels = labels.to(device)
196
197             # Forward propagation
198             outputs = model(inputs)
199             _, predicted = torch.max(outputs.data, 1)
200
201             total += labels.size(0)
202             correct += (predicted == labels).sum().item()
203
204             # Store predictions and labels
205             all_preds.extend(predicted.cpu().numpy())
206             all_labels.extend(labels.cpu().numpy())
207
208     accuracy = 100 * correct / total
209     print(f'Test Set Final Accuracy: {accuracy:.2f}%')
210
211     #
212     return accuracy, all_labels, all_preds
213
214 def plot_training_history(history, filename="training_curves.png"):
215     """
216     Plot and save the Loss and Accuracy curves during training
217     """
218     acc = history['train_acc']

```

```

219 val_acc = history['val_acc']
220 loss = history['train_loss']
221 val_loss = history['val_loss']
222
223 epochs_range = range(1, len(acc) + 1)
224
225 plt.figure(figsize=(12, 5))
226
227 # Plot Loss curves
228 plt.subplot(1, 2, 1)
229 plt.plot(epochs_range, loss, label='Training Loss')
230 plt.plot(epochs_range, val_loss, label='Validation Loss')
231 plt.legend(loc='upper right')
232 plt.title('Training and Validation Loss')
233 plt.xlabel('Epochs')
234 plt.ylabel('Loss')
235 plt.grid(True)
236
237 # Plot Accuracy curves
238 plt.subplot(1, 2, 2)
239 plt.plot(epochs_range, acc, label='Training Accuracy')
240 plt.plot(epochs_range, val_acc, label='Validation Accuracy')
241 plt.legend(loc='lower right')
242 plt.title('Training and Validation Accuracy')
243 plt.xlabel('Epochs')
244 plt.ylabel('Accuracy')
245 plt.grid(True)
246
247 plt.tight_layout()
248
249 save_path = os.path.join(FIGURE_SAVE_DIR, filename)
250 plt.savefig(save_path)
251 plt.close() # Close the figure to free memory
252 print(f"Result curves saved to: {save_path}")
253
254 def log_experiment_result(csv_file, params, metrics):
255     """
256     Append a single experiment result to a CSV file
257     :param csv_file: Path to the CSV file (e.g., 'experiment_logs.csv')
258     :param params: Dictionary containing hyperparameters (e.g., lr, batch_size, model_name)
259     :param metrics: Dictionary containing result metrics (e.g., val_acc, test_acc, time)
260     """
261     # Combine all data to be recorded
262     # Add a timestamp to distinguish different batches of experiments
263     record = {'timestamp': datetime.now().strftime("%Y-%m-%d %H:%M:%S")}
264     record.update(params) # Add hyperparameters
265     record.update(metrics) # Add results
266
267     # 2. Check if the file exists (to decide whether to write the header)
268     file_exists = os.path.isfile(csv_file)
269
270     # 3. Get all column names (Header)
271     fieldnames = list(record.keys())
272
273     # 4. Write/append data
274     try:
275         with open(csv_file, mode='a', newline='', encoding='utf-8') as f:
276             writer = csv.DictWriter(f, fieldnames=fieldnames)
277
278             # If file does not exist, write the header
279             if not file_exists:
280                 writer.writeheader()
281
282             # Write the current experiment data
283             writer.writerow(record)
284             print(f"Experiment record saved to: {csv_file}")
285
286     except Exception as e:
287         print(f"Failed to write log: {e}")
288
289 def plot_confusion_matrix(y_true, y_pred, classes, filename="confusion_matrix.png"):
290     """
291     Calculate and plot the confusion matrix
292     """

```

```

293 # Calculate confusion matrix
294 cm = confusion_matrix(y_true, y_pred)
295
296 # Create figure
297 plt.figure(figsize=(10, 8))
298
299 # Use Seaborn to draw heatmap
300 sns.heatmap(cm, annot=True, fmt='d', cmap='Blues',
301             xticklabels=classes, yticklabels=classes)
302
303 plt.title('Confusion Matrix')
304 plt.xlabel('Predicted Label')
305 plt.ylabel('Actual Label')
306 plt.xticks(rotation=45)
307 plt.yticks(rotation=45)
308 plt.tight_layout()
309
310 # Save the figure
311 save_path = os.path.join(FIGURE_SAVE_DIR, filename)
312 plt.savefig(save_path)
313 plt.close()
314 print(f"Confusion matrix saved to: {save_path}")
315
316 def run_experiment(dataset='caltech101', batch_size=32, learning_rate=0.001, num_epochs=5):
317     """
318     Main experiment function
319     """
320     print(f"\n== Running experiment: Dataset={dataset} BS={batch_size}, LR={learning_rate},
321           Epochs={num_epochs} ==")
322
323     # 1. Prepare data
324     if dataset == 'caltech101':
325         dataloaders, dataset_sizes, class_names = get_dataloaders(DATA_DIR_CALTECH101, batch_size)
326     elif dataset == 'oxford_pet':
327         dataloaders, dataset_sizes, class_names = get_dataloaders(DATA_DIR_oxford_PET, batch_size)
328     else:
329         raise ValueError(f"Unsupported dataset: {dataset}")
330     num_classes = len(class_names)
331     print(f"Detected classes: {class_names}")
332
333     # 2. Initialize model (freeze backbone)
334     model = initialize_model(num_classes, feature_extract=True)
335     model = model.to(DEVICE)
336
337     # 3. Set optimizer
338     optimizer = optim.Adam(model.fc.parameters(), lr=learning_rate)
339
340     # 4. Set loss function
341     criterion = nn.CrossEntropyLoss()
342
343     # 5. Start training
344     best_val_acc, model, history, duration = train_model(model, dataloaders, dataset_sizes,
345                                                         criterion, optimizer, num_epochs)
346
347     # 6. Save model
348     save_path = os.path.join(MODEL_SAVE_DIR, f"resnet18_{dataset}_lr{learning_rate}.pth")
349     torch.save(model.state_dict(), save_path)
350     print(f"Model saved to: {save_path}")
351
352     # 7. Plot and save training results
353     plot_filename = f"resnet18_{dataset}_lr{learning_rate}_bs{batch_size}_epochs{num_epochs}.png"
354     plot_training_history(history, plot_filename)
355
356     # 8. Load the best saved model for final testing
357     best_model_path = os.path.join(MODEL_SAVE_DIR, f"resnet18_{dataset}_lr{learning_rate}.pth")
358     model.load_state_dict(torch.load(best_model_path))
359     test_acc, y_true, y_pred = evaluate_on_test(model, dataloaders['test'], DEVICE)
360
361     # Plot confusion matrix
362     cm_filename = f"cm_cnn_{dataset}_lr{learning_rate}_bs{batch_size}.png"
363     plot_confusion_matrix(y_true, y_pred, class_names, cm_filename)
364
365     # 9. Combine all data to be recorded into a CSV file
366     # Prepare hyperparameters dictionary
367     exp_params = {

```

```

365     'dataset': dataset,
366     'model': 'resnet18',
367     'freeze_backbone': True,
368     'batch_size': batch_size,
369     'learning_rate': learning_rate,
370     'epochs': num_epochs
371 }
372
373 # Prepare results dictionary
374 exp_metrics = {
375     'best_val_acc': round(best_val_acc.item(), 4) if torch.is_tensor(best_val_acc) else round(
best_val_acc, 4),
376     'test_acc': round(test_acc, 4),
377     'training_time_sec': int(duration)
378 }
379
380 # CSV
381 log_file = LOG_CSV_FILE
382 log_experiment_result(log_file, exp_params, exp_metrics)
383
384 return history
385
386 if __name__ == "__main__":
387     # Modify parameters here for testing, including dataset selection, batch size, learning rate,
and number of epochs
388     lrs = [0.001] #0.001, 0.0005
389     batches = [64] #32, 64
390     epochs_list = [30] #15, 30
391
392     for lr in lrs:
393         for bs in batches:
394             for epochs in epochs_list:
395                 run_experiment(dataset='oxford_pet', batch_size=bs, learning_rate=lr, num_epochs=
epochs)

```

Listing 2: train_cnn.py

A.3 BoW Training Script

```

1 #train_bow.py
2 import os
3 import cv2
4 import numpy as np
5 from sklearn.cluster import MiniBatchKMeans
6 from sklearn.svm import SVC
7 from sklearn.preprocessing import StandardScaler
8 from sklearn.metrics import accuracy_score, confusion_matrix
9 import time
10 import csv
11 from datetime import datetime
12 from tqdm import tqdm
13 import matplotlib.pyplot as plt
14 import seaborn as sns
15
16 # ===== Configuration Area =====
17 # Dataset directories (Same as CNN)
18 DATA_DIR_CALTECH101 = r"D:\comp64301_cv\data\processed\caltech101"
19 DATA_DIR_oxford_pet = r"D:\comp64301_cv\data\processed\oxford_pet"
20
21 # Experiment results log file (Same file as CNN for easy comparison)
22 LOG_CSV_FILE = r"D:\comp64301_cv\results\bow_results.csv"
23 # Figure save directory
24 FIGURE_SAVE_DIR = r"D:\comp64301_cv\results\figures"
25
26 # Ensure log directory exists
27 os.makedirs(os.path.dirname(LOG_CSV_FILE), exist_ok=True)
28 os.makedirs(FIGURE_SAVE_DIR, exist_ok=True)
29 # =====
30
31 def get_file_paths(root_dir):
32     """Get file paths and labels from a directory"""
33     file_paths = []

```

```

34 labels = []
35 classes = sorted(os.listdir(root_dir))
36
37 for label_idx, class_name in enumerate(classes):
38     class_dir = os.path.join(root_dir, class_name)
39     if not os.path.isdir(class_dir):
40         continue
41     for fname in os.listdir(class_dir):
42         if fname.lower().endswith(('.jpg', '.jpeg', '.png')):
43             file_paths.append(os.path.join(class_dir, fname))
44             labels.append(label_idx)
45 return file_paths, labels, classes
46
47 def extract_sift_features(image_paths):
48     print(f"Extracting SIFT features ({len(image_paths)} images)...")
49     sift = cv2.SIFT_create()
50     descriptors_list = []
51
52     for path in tqdm(image_paths, desc="SIFT Extraction"):
53         img = cv2.imread(path, cv2.IMREAD_GRAYSCALE)
54         if img is None:
55             descriptors_list.append(None)
56             continue
57
58         _, des = sift.detectAndCompute(img, None)
59         descriptors_list.append(des)
60
61     return descriptors_list
62
63 def build_vocabulary(descriptors_list, k=100):
64     print(f"Building Visual Vocabulary (K={k})...")
65
66     valid_descriptors = [des for des in descriptors_list if des is not None]
67     if not valid_descriptors:
68         raise ValueError("No features extracted. Check image paths.")
69
70     all_descriptors = np.vstack(valid_descriptors)
71
72     kmeans = MiniBatchKMeans(n_clusters=k, batch_size=1000, random_state=42, n_init='auto')
73     kmeans.fit(all_descriptors)
74
75     return kmeans
76
77 def extract_histograms(descriptors_list, kmeans_model):
78     vocab_size = kmeans_model.n_clusters
79     histograms = []
80
81     for des in descriptors_list:
82         if des is not None:
83             # Predict which cluster each feature in this image belongs to (Visual Word)
84             words = kmeans_model.predict(des)
85             # Count the occurrences of each cluster
86             hist, _ = np.histogram(words, bins=range(vocab_size + 1))
87             # Normalize (L2 Norm) - This step is important for SVM
88             norm = np.linalg.norm(hist)
89             if norm == 0:
90                 histograms.append(hist)
91             else:
92                 histograms.append(hist / norm)
93         else:
94             # If no features, give a zero vector
95             histograms.append(np.zeros(vocab_size))
96
97     return np.array(histograms)
98
99 def log_experiment_result(csv_file, params, metrics):
100     record = {'timestamp': datetime.now().strftime("%Y-%m-%d %H:%M:%S")}
101     record.update(params)
102     record.update(metrics)
103
104     file_exists = os.path.isfile(csv_file)
105     fieldnames = list(record.keys())
106
107     try:

```

```

108         with open(csv_file, mode='a', newline='', encoding='utf-8') as f:
109             writer = csv.DictWriter(f, fieldnames=fieldnames)
110             if not file_exists:
111                 writer.writeheader()
112                 writer.writerow(record)
113                 print(f"Experiment record saved to: {csv_file}")
114     except Exception as e:
115         print(f"Failed to write log: {e}")
116
117 def run_experiment(dataset='caltech101', vocab_size=100, svm_c=1.0, svm_kernel='rbf'):
118     """
119     Main experiment function for BoW
120     """
121     start_time = time.time()
122     print(f"\n=== Running BoW Experiment: Dataset={dataset}, K={vocab_size}, SVM_C={svm_c},
123           SVM_Kernel={svm_kernel} ===")
124
125     # 1. Select Dataset Path
126     if dataset == 'caltech101':
127         data_dir = DATA_DIR_CALTECH101
128     elif dataset == 'oxford_pet':
129         data_dir = DATA_DIR_OXFORD_PET
130     else:
131         raise ValueError(f"Unsupported dataset: {dataset}")
132
133     train_dir = os.path.join(data_dir, "train")
134     test_dir = os.path.join(data_dir, "test")
135
136     train_paths, train_y, classes = get_file_paths(train_dir)
137     test_paths, test_y, _ = get_file_paths(test_dir)
138
139     print(f"Classes: {len(classes)}")
140
141     # 2. Feature Extraction
142     train_des = extract_sift_features(train_paths)
143     test_des = extract_sift_features(test_paths)
144
145     # 3. Clustering (Vocabulary Building)
146     kmeans = build_vocabulary(train_des, k=vocab_size)
147
148     # 4. Encoding
149     print("Generating histograms...")
150     X_train = extract_histograms(train_des, kmeans)
151     X_test = extract_histograms(test_des, kmeans)
152
153     # 5. Standardization
154     scaler = StandardScaler()
155     X_train = scaler.fit_transform(X_train)
156     X_test = scaler.transform(X_test)
157
158     # 6. SVM Training
159     print(f"Training SVM (C={svm_c}, Kernel={svm_kernel})...")
160     clf = SVC(kernel=svm_kernel, C=svm_c, random_state=42)
161     clf.fit(X_train, train_y)
162
163     # 7. Evaluation
164     print("Evaluating on Test set...")
165     y_pred = clf.predict(X_test)
166     test_acc = accuracy_score(test_y, y_pred)
167
168     # 8. Plot Confusion Matrix
169     cm = confusion_matrix(test_y, y_pred)
170     plt.figure(figsize=(10, 8))
171     sns.heatmap(cm, annot=True, fmt='d', cmap='Greens',
172                xticklabels=classes, yticklabels=classes)
173     plt.title("Confusion Matrix")
174     plt.xlabel("Predicted Label")
175     plt.ylabel("Actual Label")
176     plt.xticks(rotation=45)
177     plt.yticks(rotation=45)
178     plt.tight_layout()
179     plt.savefig(os.path.join(FIGURE_SAVE_DIR, f"cm_bow_{dataset}_K{vocab_size}_C{svm_c}_Kernel{
180           svm_kernel}.png"))
181     plt.close()

```

```

180     duration = time.time() - start_time
181
182
183     print("-" * 30)
184     print(f"BoW Result: Test Acc = {test_acc*100:.2f}%, Time = {duration:.1f}s")
185     print("-" * 30)
186
187     # 9. Log Results
188     exp_params = {
189         'dataset': dataset,
190         'method': 'BoW+SVM',
191         'vocab_size': vocab_size,
192         'svm_c': svm_c,
193         'svm_kernel': svm_kernel
194     }
195
196     exp_metrics = {
197         'test_acc': round(test_acc, 4),
198         'training_time_sec': int(duration)
199     }
200
201     log_experiment_result(LOG_CSV_FILE, exp_params, exp_metrics)
202
203 if __name__ == "__main__":
204     datasets = ['caltech101', 'oxford_pet']# 'caltech101', 'oxford_pet'
205     vocab_sizes = [100]# 50, 100, 200, 500
206     svm_cs = [1.0]# 0.1, 1.0, 10.0
207     svm_kernels = ['rbf']# 'linear', 'rbf'
208
209     for dataset in datasets:
210         for vocab_size in vocab_sizes:
211             for svm_c in svm_cs:
212                 for svm_kernel in svm_kernels:
213                     run_experiment(dataset=dataset, vocab_size=vocab_size, svm_c=svm_c, svm_kernel
=svm_kernel)

```

Listing 3: train_bow.py