

```
from google.colab import drive
drive.mount('/content/drive')
```

Drive already mounted at /content/drive; to attempt to forcibly remount, call drive.mount("/content/drive", force_remount=True)

```
import os
import shutil
import math
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
import seaborn as sns

import tensorflow as tf
from tensorflow.keras.models import Sequential, Model
from tensorflow.keras.layers import (
    Conv2D, MaxPooling2D, Flatten,
    Dense, Dropout, BatchNormalization,
    GlobalAveragePooling2D
)
from tensorflow.keras.preprocessing.image import ImageDataGenerator
from tensorflow.keras.callbacks import EarlyStopping, ModelCheckpoint, ReduceLROnPlateau
from tensorflow.keras.optimizers import Adam

from tensorflow.keras.applications import MobileNetV2, ResNet50, EfficientNetB0
from tensorflow.keras.applications.mobilenet_v2 import preprocess_input as mobilenet_preprocess
from tensorflow.keras.applications.resnet50 import preprocess_input as resnet_preprocess
from tensorflow.keras.applications.efficientnet import preprocess_input as efficientnet_preprocess

from sklearn.metrics import classification_report, confusion_matrix

print("TensorFlow Version:", tf.__version__)
```

TensorFlow Version: 2.20.0

```
DATASET_PATH = r"/content/drive/MyDrive/Final Project AI/asl_dataset"

VALID_CLASSES = [
    '0','1','2','3','4','5','6','7','8','9',
    'a','b','c','d','e','f','g','h','i','j','k','l','m',
    'n','o','p','q','r','s','t','u','v','w','x','y','z'
]

NUM_CLASSES = len(VALID_CLASSES)

print("Dataset Path:", DATASET_PATH)
print("Number of Classes:", NUM_CLASSES)

if not os.path.isdir(DATASET_PATH):
    raise Exception("Dataset path is not correct. Please check DATASET_PATH.")
else:
    print("Dataset folder found successfully.")
```

Dataset Path: /content/drive/MyDrive/Final Project AI/asl_dataset
 Number of Classes: 36
 Dataset folder found successfully.

```
# =====
# Part 4: Create Clean Dataset
# =====

SRC = DATASET_PATH
DEST = r"/content/drive/MyDrive/clean_dataset"

if os.path.exists(DEST):
    shutil.rmtree(DEST)

os.makedirs(DEST)

print("Copying dataset...\n")

total = 0

for cls in VALID_CLASSES:
    src_cls = os.path.join(SRC, cls)
    dest_cls = os.path.join(DEST, cls)

    os.makedirs(dest_cls, exist_ok=True)
```

```

if not os.path.exists(src_cls):
    print(f"Missing folder: {src_cls}")
    continue

imgs = [
    f for f in os.listdir(src_cls)
    if f.lower().endswith(('.jpg', '.jpeg', '.png'))
    and os.path.isfile(os.path.join(src_cls, f))
]

for img in imgs:
    shutil.copy2(os.path.join(src_cls, img), os.path.join(dest_cls, img))

total += len(imgs)
print(f"{cls}: {len(imgs)} images copied")

print(f"\nDone. Total images copied: {total}")
print(f"Clean dataset saved at: {DEST}")

DATASET_PATH = DEST

```

Copying dataset...

```

0: 70 images copied
1: 70 images copied
2: 70 images copied
3: 70 images copied
4: 70 images copied
5: 70 images copied
6: 70 images copied
7: 70 images copied
8: 70 images copied
9: 70 images copied
a: 71 images copied
b: 70 images copied
c: 70 images copied
d: 70 images copied
e: 70 images copied
f: 70 images copied
g: 70 images copied
h: 70 images copied
i: 70 images copied
j: 70 images copied
k: 70 images copied
l: 70 images copied
m: 70 images copied
n: 70 images copied
o: 70 images copied
p: 70 images copied
q: 70 images copied
r: 70 images copied
s: 70 images copied
t: 65 images copied
u: 70 images copied
v: 70 images copied
w: 70 images copied
x: 70 images copied
y: 70 images copied
z: 70 images copied

```

Done. Total images copied: 2516
Clean dataset saved at: /content/drive/MyDrive/clean_dataset

```

# =====
# Part 5: Visualize Sample Images
# =====

from PIL import Image

cols = 7
rows = math.ceil(NUM_CLASSES / cols)

fig, axes = plt.subplots(rows, cols, figsize=(20, rows * 3))
fig.suptitle("ASL Dataset - Sample Images per Class", fontsize=16, fontweight="bold")

axes = axes.flatten()

for idx, cls in enumerate(VALID_CLASSES):
    cls_path = os.path.join(DATASET_PATH, cls)

    img_files = [
        f for f in os.listdir(cls_path)
        if f.lower().endswith(('.jpg', '.jpeg', '.png'))
    ]

```

```

if not img_files:
    axes[idx].axis("off")
    continue

img_path = os.path.join(cls_path, img_files[0])
img = Image.open(img_path)

axes[idx].imshow(img)
axes[idx].set_title(cls.upper(), fontsize=12, fontweight="bold")
axes[idx].axis("off")

for idx in range(len(VALID_CLASSES), len(axes)):
    axes[idx].axis("off")

plt.tight_layout()
plt.savefig("sample_images.png", dpi=150, bbox_inches="tight")
plt.show()

print(f"Visualized {NUM_CLASSES} classes.")

```



Visualized 36 classes.

```

# =====
# Part 6: Define Hyperparameters

```

```
# =====

# Custom CNN Settings
CNN_IMG_SIZE = 64
CNN_BATCH_SIZE = 16
CNN_EPOCHS = 50
CNN_LR = 0.001

# Transfer Learning Settings
TL_IMG_SIZE = 224
TL_BATCH_SIZE = 16
TL_EPOCHS = 15
TL_LR = 0.0001

print("CNN Image Size:", CNN_IMG_SIZE)
print("Transfer Learning Image Size:", TL_IMG_SIZE)
print("Number of Classes:", NUM_CLASSES)
```

```
CNN Image Size: 64
Transfer Learning Image Size: 224
Number of Classes: 36
```

```
# =====
# Part 7: Prepare Data Generators for Custom CNN
# لنموذج CNN تجهيز البيانات والـ Data Augmentation
# =====

cnn_train_datagen = ImageDataGenerator(
    rescale=1./255,
    validation_split=0.2,
    rotation_range=20,
    width_shift_range=0.15,
    height_shift_range=0.15,
    zoom_range=0.15,
    shear_range=0.1,
    brightness_range=[0.7, 1.3],
    horizontal_flip=False,
    fill_mode="nearest"
)

cnn_val_datagen = ImageDataGenerator(
    rescale=1./255,
    validation_split=0.2
)

cnn_train_generator = cnn_train_datagen.flow_from_directory(
    DATASET_PATH,
    target_size=(CNN_IMG_SIZE, CNN_IMG_SIZE),
    batch_size=CNN_BATCH_SIZE,
    class_mode="categorical",
    subset="training",
    shuffle=True,
    seed=42
)

cnn_val_generator = cnn_val_datagen.flow_from_directory(
    DATASET_PATH,
    target_size=(CNN_IMG_SIZE, CNN_IMG_SIZE),
    batch_size=CNN_BATCH_SIZE,
    class_mode="categorical",
    subset="validation",
    shuffle=False,
    seed=42
)

CLASS_NAMES = list(cnn_train_generator.class_indices.keys())

print("Classes used:", CLASS_NAMES)
print("Training images:", cnn_train_generator.n)
print("Validation images:", cnn_val_generator.n)
print("Train batches:", len(cnn_train_generator))
print("Validation batches:", len(cnn_val_generator))
```

```
Found 2013 images belonging to 36 classes.
Found 503 images belonging to 36 classes.
Classes used: ['0', '1', '2', '3', '4', '5', '6', '7', '8', '9', 'a', 'b', 'c', 'd', 'e', 'f', 'g', 'h', 'i', 'j', 'k', 'l',
Training images: 2013
Validation images: 503
Train batches: 126
Validation batches: 32
```

```
# =====
# Part 8: Build Custom CNN Model
# =====

cnn_model = Sequential([
    Conv2D(32, (3, 3), activation="relu", input_shape=(CNN_IMG_SIZE, CNN_IMG_SIZE, 3)),
    BatchNormalization(),
    MaxPooling2D(2, 2),

    Conv2D(64, (3, 3), activation="relu"),
    BatchNormalization(),
    MaxPooling2D(2, 2),

    Conv2D(128, (3, 3), activation="relu"),
    BatchNormalization(),
    MaxPooling2D(2, 2),

    Conv2D(256, (3, 3), activation="relu"),
    BatchNormalization(),
    MaxPooling2D(2, 2),

    Flatten(),

    Dense(256, activation="relu"),
    Dropout(0.5),

    Dense(NUM_CLASSES, activation="softmax")
])

cnn_model.compile(
    optimizer=Adam(learning_rate=CNN_LR),
    loss="categorical_crossentropy",
    metrics=["accuracy"]
)

cnn_model.summary()
```

/usr/local/lib/python3.12/dist-packages/keras/src/layers/convolutional/base_conv.py:113: UserWarning: Do not pass an `input\_shape` argument to `Conv2D` or `Conv3D` layers. The input shape is inferred from the first batch data. In the future, this warning will become an error.

Model: "sequential"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 62, 62, 32)	896
batch_normalization (BatchNormalization)	(None, 62, 62, 32)	128
max_pooling2d (MaxPooling2D)	(None, 31, 31, 32)	0
conv2d_1 (Conv2D)	(None, 29, 29, 64)	18,496
batch_normalization_1 (BatchNormalization)	(None, 29, 29, 64)	256
max_pooling2d_1 (MaxPooling2D)	(None, 14, 14, 64)	0
conv2d_2 (Conv2D)	(None, 12, 12, 128)	73,856
batch_normalization_2 (BatchNormalization)	(None, 12, 12, 128)	512
max_pooling2d_2 (MaxPooling2D)	(None, 6, 6, 128)	0
conv2d_3 (Conv2D)	(None, 4, 4, 256)	295,168
batch_normalization_3 (BatchNormalization)	(None, 4, 4, 256)	1,024
max_pooling2d_3 (MaxPooling2D)	(None, 2, 2, 256)	0
flatten (Flatten)	(None, 1024)	0
dense (Dense)	(None, 256)	262,400
dropout (Dropout)	(None, 256)	0
dense_1 (Dense)	(None, 36)	9,252

Total params: 661,988 (2.53 MB)
 Trainable params: 661,028 (2.52 MB)
 Non-trainable params: 960 (3.75 KB)

```
# =====
# Part 9: Train Custom CNN Model
# الأساسي تدريب نموذج CNN
# =====
```

```

cnn_callbacks = [
    EarlyStopping(
        monitor="val_loss",
        patience=7,
        restore_best_weights=True
    ),

    ReduceLROnPlateau(
        monitor="val_loss",
        factor=0.3,
        patience=4,
        min_lr=1e-6
    ),

    ModelCheckpoint(
        "custom_cnn_model.h5",
        monitor="val_accuracy",
        save_best_only=True,
        verbose=1
    )
]

cnn_history = cnn_model.fit(
    cnn_train_generator,
    validation_data=cnn_val_generator,
    epochs=CNN_EPOCHS,
    callbacks=cnn_callbacks
)

cnn_loss, cnn_acc = cnn_model.evaluate(
    cnn_val_generator,
    verbose=0
)

print(f"Custom CNN Validation Accuracy: {cnn_acc * 100:.2f}%")
print(f"Custom CNN Validation Loss      : {cnn_loss:.4f}")

```

```

Epoch 1/50
126/126 ————— 0s 266ms/step - accuracy: 0.0855 - loss: 4.0288
Epoch 1: val_accuracy improved from None to 0.05567, saving model to custom_cnn_model.h5
WARNING:absl:You are saving your model as an HDF5 file via `model.save()` or `keras.saving.save_model(model)`. This file f

Epoch 1: finished saving model to custom_cnn_model.h5
126/126 ————— 42s 295ms/step - accuracy: 0.1465 - loss: 3.3639 - val_accuracy: 0.0557 - val_loss: 5.2991 -
Epoch 2/50
126/126 ————— 0s 261ms/step - accuracy: 0.2899 - loss: 2.4516
Epoch 2: val_accuracy improved from 0.05567 to 0.08350, saving model to custom_cnn_model.h5
WARNING:absl:You are saving your model as an HDF5 file via `model.save()` or `keras.saving.save_model(model)`. This file f

Epoch 2: finished saving model to custom_cnn_model.h5
126/126 ————— 38s 303ms/step - accuracy: 0.3135 - loss: 2.3226 - val_accuracy: 0.0835 - val_loss: 3.8478 -
Epoch 3/50
126/126 ————— 0s 259ms/step - accuracy: 0.4003 - loss: 1.9357
Epoch 3: val_accuracy improved from 0.08350 to 0.34791, saving model to custom_cnn_model.h5
WARNING:absl:You are saving your model as an HDF5 file via `model.save()` or `keras.saving.save_model(model)`. This file f

Epoch 3: finished saving model to custom_cnn_model.h5
126/126 ————— 36s 283ms/step - accuracy: 0.4069 - loss: 1.8968 - val_accuracy: 0.3479 - val_loss: 1.9724 -
Epoch 4/50
126/126 ————— 0s 270ms/step - accuracy: 0.4711 - loss: 1.6611
Epoch 4: val_accuracy improved from 0.34791 to 0.80119, saving model to custom_cnn_model.h5
WARNING:absl:You are saving your model as an HDF5 file via `model.save()` or `keras.saving.save_model(model)`. This file f

Epoch 4: finished saving model to custom_cnn_model.h5
126/126 ————— 43s 296ms/step - accuracy: 0.4814 - loss: 1.6038 - val_accuracy: 0.8012 - val_loss: 0.5947 -
Epoch 5/50
126/126 ————— 0s 259ms/step - accuracy: 0.5317 - loss: 1.3819
Epoch 5: val_accuracy did not improve from 0.80119
126/126 ————— 36s 282ms/step - accuracy: 0.5405 - loss: 1.3289 - val_accuracy: 0.7813 - val_loss: 0.6546 -
Epoch 6/50
126/126 ————— 0s 262ms/step - accuracy: 0.6035 - loss: 1.2347
Epoch 6: val_accuracy did not improve from 0.80119
126/126 ————— 38s 298ms/step - accuracy: 0.5971 - loss: 1.1900 - val_accuracy: 0.7853 - val_loss: 0.4938 -
Epoch 7/50
126/126 ————— 0s 261ms/step - accuracy: 0.6731 - loss: 0.9523
Epoch 7: val_accuracy improved from 0.80119 to 0.83300, saving model to custom_cnn_model.h5
WARNING:absl:You are saving your model as an HDF5 file via `model.save()` or `keras.saving.save_model(model)`. This file f

Epoch 7: finished saving model to custom_cnn_model.h5
126/126 ————— 36s 286ms/step - accuracy: 0.6612 - loss: 0.9861 - val_accuracy: 0.8330 - val_loss: 0.4678 -
Epoch 8/50
126/126 ————— 0s 262ms/step - accuracy: 0.6470 - loss: 0.9694
Epoch 8: val_accuracy did not improve from 0.83300
126/126 ————— 37s 292ms/step - accuracy: 0.6642 - loss: 0.9340 - val_accuracy: 0.7435 - val_loss: 0.7233 -
Epoch 9/50

```

```

126/126 0s 267ms/step - accuracy: 0.7215 - loss: 0.8233
Epoch 9: val_accuracy did not improve from 0.83300
126/126 37s 291ms/step - accuracy: 0.7119 - loss: 0.8372 - val_accuracy: 0.8151 - val_loss: 0.5146 -
Epoch 10/50
126/126 0s 258ms/step - accuracy: 0.7037 - loss: 0.7904
Epoch 10: val_accuracy did not improve from 0.83300
126/126 36s 282ms/step - accuracy: 0.7248 - loss: 0.7743 - val_accuracy: 0.7932 - val_loss: 0.5255 -
Epoch 11/50
126/126 0s 260ms/step - accuracy: 0.7000 - loss: 0.8550

```

```

# =====
# Part 10: Plot Training Curves Function
# =====

def plot_training_curves(history, model_name):
    plt.figure(figsize=(14, 5))

    plt.subplot(1, 2, 1)
    plt.plot(history.history["accuracy"], label="Train")
    plt.plot(history.history["val_accuracy"], label="Validation")
    plt.title(f"{model_name} Accuracy")
    plt.xlabel("Epoch")
    plt.ylabel("Accuracy")
    plt.legend()

    plt.subplot(1, 2, 2)
    plt.plot(history.history["loss"], label="Train")
    plt.plot(history.history["val_loss"], label="Validation")
    plt.title(f"{model_name} Loss")
    plt.xlabel("Epoch")
    plt.ylabel("Loss")
    plt.legend()

    plt.tight_layout()
    plt.show()


```

```

# =====
# Part 11: Plot CNN Training Curves
# =====

plot_training_curves(cnn_history, "Custom CNN")

```



```

# =====
# Part 12: Evaluate Custom CNN Model
# =====

cnn_val_generator.reset()

cnn_preds = cnn_model.predict(cnn_val_generator)
cnn_y_pred = np.argmax(cnn_preds, axis=1)
cnn_y_true = cnn_val_generator.classes

print("Custom CNN Classification Report")
print(classification_report(
    cnn_y_true,
    cnn_y_pred,
    target_names=CLASS_NAMES
))

```

```
cnn_cm = confusion_matrix(cnn_y_true, cnn_y_pred)

plt.figure(figsize=(14, 12))
sns.heatmap(
    cnn_cm,
    xticklabels=CLASS_NAMES,
    yticklabels=CLASS_NAMES,
    cmap="Blues",
    fmt="d"
)
plt.title("Custom CNN Confusion Matrix")
plt.xlabel("Predicted")
plt.ylabel("True")
plt.show()
```



```

32/32 ----- 3s 100ms/step
Custom CNN Classification Report
      precision    recall  f1-score   support

     0         0.67      0.29      0.40         14
     1         0.82      1.00      0.90         14
     2         0.88      1.00      0.93         14
     3         1.00      1.00      1.00         14
     4         1.00      1.00      1.00         14
     5         1.00      1.00      1.00         14
     6         1.00      0.50      0.67         14
     7         1.00      1.00      1.00         14
     8         1.00      1.00      1.00         14
     9         1.00      1.00      1.00         14
    a         1.00      1.00      1.00         14

```

```

# =====
# Part 13: Create Transfer Learning Data Generators
# MobileNetV2 / ResNet50 / EfficientNetB0
# =====

def create_transfer_generators(preprocess_function):
    train_datagen_tl = ImageDataGenerator(
        preprocessing_function=preprocess_function,
        validation_split=0.2,
        rotation_range=20,
        width_shift_range=0.15,
        height_shift_range=0.15,
        zoom_range=0.15,
        shear_range=0.1,
        brightness_range=[0.7, 1.3],
        horizontal_flip=False,
        fill_mode="nearest"
    )

    val_datagen_tl = ImageDataGenerator(
        preprocessing_function=preprocess_function,
        validation_split=0.2
    )

    train_gen = train_datagen_tl.flow_from_directory(
        DATASET_PATH,
        target_size=(TL_IMG_SIZE, TL_IMG_SIZE),
        batch_size=TL_BATCH_SIZE,
        class_mode="categorical",
        subset="training",
        shuffle=True,
        seed=42
    )

    val_gen = val_datagen_tl.flow_from_directory(
        DATASET_PATH,
        target_size=(TL_IMG_SIZE, TL_IMG_SIZE),
        batch_size=TL_BATCH_SIZE,
        class_mode="categorical",
        subset="validation",
        shuffle=False,
        seed=42
    )

    return train_gen, val_gen

```

```

# =====
# Part 14: Build Transfer Learning Model Function
# إنشاء دالة لبناء نموذج Transfer Learning
# =====

def build_transfer_learning_model(base_model, model_name):
    base_model.trainable = False

    x = base_model.output
    x = GlobalAveragePooling2D()(x)
    x = Dense(256, activation="relu")(x)
    x = Dropout(0.5)(x)

    output = Dense(NUM_CLASSES, activation="softmax")(x)

    model = Model(
        inputs=base_model.input,
        outputs=output,
        name=model_name
    )

```

```
model.compile(
    optimizer=Adam(learning_rate=TL_LR),
    loss="categorical_crossentropy",
    metrics=["accuracy"]
)

return model
```

```
# =====
# Part 15: Train MobileNetV2 Model
# باستخدام MobileNetV2 لتدريب نموذج Transfer Learning
# =====

mobilenet_train_gen, mobilenet_val_gen = create_transfer_generators(mobilenet_preprocess)

mobilenet_base = MobileNetV2(
    weights="imagenet",
    include_top=False,
    input_shape=(TL_IMG_SIZE, TL_IMG_SIZE, 3)
)

mobilenet_model = build_transfer_learning_model(
    mobilenet_base,
    "MobileNetV2_ASL"
)

mobilenet_model.summary()

mobilenet_history = mobilenet_model.fit(
    mobilenet_train_gen,
    validation_data=mobilenet_val_gen,
    epochs=TL_EPOCHS
)

mobilenet_loss, mobilenet_acc = mobilenet_model.evaluate(
    mobilenet_val_gen,
    verbose=0
)

print(f"MobileNetV2 Validation Accuracy: {mobilenet_acc * 100:.2f}%")
print(f"MobileNetV2 Validation Loss : {mobilenet_loss:.4f}")

plot_training_curves(mobilenet_history, "MobileNetV2")
```



```
Found 503 images belonging to 36 classes.
```

Model: "MobileNetV2 ASL"

Model: "MobileNetV2 ASL"

Layer (type)	Output Shape	Param #	Connected to
input_layer_1 (InputLayer)	(None, 224, 224, 3)	0	-
Conv1 (Conv2D)	(None, 112, 112, 32)	864	input_layer_1[0]...
bn_Conv1 (BatchNormalizatio...	(None, 112, 112, 32)	128	Conv1[0][0]
Conv1_relu (ReLU)	(None, 112, 112, 32)	0	bn_Conv1[0][0]
expanded_conv_dept... (DepthwiseConv2D)	(None, 112, 112, 32)	288	Conv1_relu[0][0]
expanded_conv_dept... (BatchNormalizatio...	(None, 112, 112, 32)	128	expanded_conv_de...
expanded_conv_dept... (ReLU)	(None, 112, 112, 32)	0	expanded_conv_de...
expanded_conv_proj... (Conv2D)	(None, 112, 112, 16)	512	expanded_conv_de...
expanded_conv_proj... (BatchNormalizatio...	(None, 112, 112, 16)	64	expanded_conv_pr...
block_1_expand (Conv2D)	(None, 112, 112, 96)	1,536	expanded_conv_pr...
block_1_expand_BN (BatchNormalizatio...	(None, 112, 112, 96)	384	block_1_expand[0...
block_1_expand_relu (ReLU)	(None, 112, 112, 96)	0	block_1_expand_B...
block_1_pad (ZeroPadding2D)	(None, 113, 113, 96)	0	block_1_expand_r...
block_1_depthwise (DepthwiseConv2D)	(None, 56, 56, 96)	864	block_1_pad[0][0]
block_1_depthwise_... (BatchNormalizatio...	(None, 56, 56, 96)	384	block_1_depthwis...
block_1_depthwise_... (ReLU)	(None, 56, 56, 96)	0	block_1_depthwis...
block_1_project (Conv2D)	(None, 56, 56, 24)	2,304	block_1_depthwis...
block_1_project_BN (BatchNormalizatio...	(None, 56, 56, 24)	96	block_1_project[...
block_2_expand (Conv2D)	(None, 56, 56, 144)	3,456	block_1_project_...
block_2_expand_BN (BatchNormalizatio...	(None, 56, 56, 144)	576	block_2_expand[0...
block_2_expand_relu (ReLU)	(None, 56, 56, 144)	0	block_2_expand_B...
block_2_depthwise (DepthwiseConv2D)	(None, 56, 56, 144)	1,296	block_2_expand_r...
block_2_depthwise_... (BatchNormalizatio...	(None, 56, 56, 144)	576	block_2_depthwis...
block_2_depthwise_... (ReLU)	(None, 56, 56, 144)	0	block_2_depthwis...
block_2_project (Conv2D)	(None, 56, 56, 24)	3,456	block_2_depthwis...
block_2_project_BN (BatchNormalizatio...	(None, 56, 56, 24)	96	block_2_project[...
block_2_add (Add)	(None, 56, 56, 24)	0	block_1_project_... block_2_project_...
block_3_expand (Conv2D)	(None, 56, 56, 144)	3,456	block_2_add[0][0]
block_3_expand_BN (BatchNormalizatio...	(None, 56, 56, 144)	576	block_3_expand[0...
block_3_expand_relu (ReLU)	(None, 56, 56, 144)	0	block_3_expand_B...

block_3_pad (ZeroPadding2D)	(None, 57, 57, 144)	0	block_3_expand_r...
block_3_depthwise (DepthwiseConv2D)	(None, 28, 28, 144)	1,296	block_3_pad[0][0]
block_3_depthwise_... (BatchNormalizatio...	(None, 28, 28, 144)	576	block_3_depthwis...
block_3_depthwise_... (ReLU)	(None, 28, 28, 144)	0	block_3_depthwis...
block_3_project (Conv2D)	(None, 28, 28, 32)	4,608	block_3_depthwis...
block_3_project_BN (BatchNormalizatio...	(None, 28, 28, 32)	128	block_3_project[...
block_4_expand (Conv2D)	(None, 28, 28, 192)	6,144	block_3_project_...
block_4_expand_BN (BatchNormalizatio...	(None, 28, 28, 192)	768	block_4_expand[0...
block_4_expand_relu (ReLU)	(None, 28, 28, 192)	0	block_4_expand_B...
block_4_depthwise (DepthwiseConv2D)	(None, 28, 28, 192)	1,728	block_4_expand_r...
block_4_depthwise_... (BatchNormalizatio...	(None, 28, 28, 192)	768	block_4_depthwis...
block_4_depthwise_... (ReLU)	(None, 28, 28, 192)	0	block_4_depthwis...
block_4_project (Conv2D)	(None, 28, 28, 32)	6,144	block_4_depthwis...
block_4_project_BN (BatchNormalizatio...	(None, 28, 28, 32)	128	block_4_project[...
block_4_add (Add)	(None, 28, 28, 32)	0	block_3_project_... block_4_project_...
block_5_expand (Conv2D)	(None, 28, 28, 192)	6,144	block_4_add[0][0]
block_5_expand_BN (BatchNormalizatio...	(None, 28, 28, 192)	768	block_5_expand[0...
block_5_expand_relu (ReLU)	(None, 28, 28, 192)	0	block_5_expand_B...
block_5_depthwise (DepthwiseConv2D)	(None, 28, 28, 192)	1,728	block_5_expand_r...
block_5_depthwise_... (BatchNormalizatio...	(None, 28, 28, 192)	768	block_5_depthwis...
block_5_depthwise_... (ReLU)	(None, 28, 28, 192)	0	block_5_depthwis...
block_5_project (Conv2D)	(None, 28, 28, 32)	6,144	block_5_depthwis...
block_5_project_BN (BatchNormalizatio...	(None, 28, 28, 32)	128	block_5_project[...
block_5_add (Add)	(None, 28, 28, 32)	0	block_4_add[0][0... block_5_project_...
block_6_expand (Conv2D)	(None, 28, 28, 192)	6,144	block_5_add[0][0]
block_6_expand_BN (BatchNormalizatio...	(None, 28, 28, 192)	768	block_6_expand[0...
block_6_expand_relu (ReLU)	(None, 28, 28, 192)	0	block_6_expand_B...
block_6_pad (ZeroPadding2D)	(None, 29, 29, 192)	0	block_6_expand_r...
block_6_depthwise (DepthwiseConv2D)	(None, 14, 14, 192)	1,728	block_6_pad[0][0]
block_6_depthwise_... (BatchNormalizatio...	(None, 14, 14, 192)	768	block_6_depthwis...
block_6_depthwise_... (ReLU)	(None, 14, 14, 192)	0	block_6_depthwis...
block_6_project (Conv2D)	(None, 14, 14, 64)	12,288	block_6_depthwis...
block_6_project_BN (BatchNormalizatio...	(None, 14, 14, 64)	256	block_6_project[...

(BatchNormalization)	(None, 14, 14, 384)	24,576	block_6_project_...
block_7_expand (Conv2D)	(None, 14, 14, 384)	1,536	block_7_expand[0...
block_7_expand_BN (BatchNormalization)	(None, 14, 14, 384)	0	block_7_expand_B...
block_7_expand_relu (ReLU)	(None, 14, 14, 384)	3,456	block_7_expand_r...
block_7_depthwise (DepthwiseConv2D)	(None, 14, 14, 384)	1,536	block_7_depthwis...
block_7_depthwise_... (BatchNormalization)	(None, 14, 14, 384)	0	block_7_depthwis...
block_7_depthwise_... (ReLU)	(None, 14, 14, 384)	24,576	block_7_depthwis...
block_7_project (Conv2D)	(None, 14, 14, 64)	256	block_7_project[...
block_7_project_BN (BatchNormalization)	(None, 14, 14, 64)	0	block_6_project_... block_7_project_...
block_7_add (Add)	(None, 14, 14, 64)	24,576	block_7_add[0][0]
block_8_expand (Conv2D)	(None, 14, 14, 384)	1,536	block_8_expand[0...
block_8_expand_BN (BatchNormalization)	(None, 14, 14, 384)	0	block_8_expand_B...
block_8_expand_relu (ReLU)	(None, 14, 14, 384)	3,456	block_8_expand_r...
block_8_depthwise (DepthwiseConv2D)	(None, 14, 14, 384)	1,536	block_8_depthwis...
block_8_depthwise_... (BatchNormalization)	(None, 14, 14, 384)	0	block_8_depthwis...
block_8_depthwise_... (ReLU)	(None, 14, 14, 384)	24,576	block_8_depthwis...
block_8_project (Conv2D)	(None, 14, 14, 64)	256	block_8_project[...
block_8_project_BN (BatchNormalization)	(None, 14, 14, 64)	0	block_7_add[0][0... block_8_project_...
block_8_add (Add)	(None, 14, 14, 64)	24,576	block_8_add[0][0]
block_9_expand (Conv2D)	(None, 14, 14, 384)	1,536	block_9_expand[0...
block_9_expand_BN (BatchNormalization)	(None, 14, 14, 384)	0	block_9_expand_B...
block_9_expand_relu (ReLU)	(None, 14, 14, 384)	3,456	block_9_expand_r...
block_9_depthwise (DepthwiseConv2D)	(None, 14, 14, 384)	1,536	block_9_depthwis...
block_9_depthwise_... (BatchNormalization)	(None, 14, 14, 384)	0	block_9_depthwis...
block_9_depthwise_... (ReLU)	(None, 14, 14, 384)	24,576	block_9_depthwis...
block_9_project (Conv2D)	(None, 14, 14, 64)	256	block_9_project[...
block_9_project_BN (BatchNormalization)	(None, 14, 14, 64)	0	block_8_add[0][0... block_9_project_...
block_9_add (Add)	(None, 14, 14, 64)	24,576	block_9_add[0][0]
block_10_expand (Conv2D)	(None, 14, 14, 384)	1,536	block_10_expand[...
block_10_expand_BN (BatchNormalization)	(None, 14, 14, 384)	0	block_10_expand_...
block_10_expand_re... (ReLU)	(None, 14, 14, 384)	3,456	block_10_expand_...
block_10_depthwise (DepthwiseConv2D)	(None, 14, 14, 384)	1,536	block_10_depthwi...
block_10_depthwise... (BatchNormalization)	(None, 14, 14, 384)	0	block_10_depthwi...
block_10_depthwise... (ReLU)	(None, 14, 14, 384)	0	block_10_depthwi...

block_10_depthwise (ReLU)	(None, 14, 14, 384)	36,864	block_10_depthwi...
block_10_project (Conv2D)	(None, 14, 14, 96)	36,864	block_10_depthwi...
block_10_project_BN (BatchNormalizatio...	(None, 14, 14, 96)	384	block_10_project...
block_11_expand (Conv2D)	(None, 14, 14, 576)	55,296	block_10_project...
block_11_expand_BN (BatchNormalizatio...	(None, 14, 14, 576)	2,304	block_11_expand[...
block_11_expand_re... (ReLU)	(None, 14, 14, 576)	0	block_11_expand_...
block_11_depthwise (DepthwiseConv2D)	(None, 14, 14, 576)	5,184	block_11_expand_...
block_11_depthwise... (BatchNormalizatio...	(None, 14, 14, 576)	2,304	block_11_depthwi...
block_11_depthwise... (ReLU)	(None, 14, 14, 576)	0	block_11_depthwi...
block_11_project (Conv2D)	(None, 14, 14, 96)	55,296	block_11_depthwi...
block_11_project_BN (BatchNormalizatio...	(None, 14, 14, 96)	384	block_11_project...
block_11_add (Add)	(None, 14, 14, 96)	0	block_10_project... block_11_project...
block_12_expand (Conv2D)	(None, 14, 14, 576)	55,296	block_11_add[0][...
block_12_expand_BN (BatchNormalizatio...	(None, 14, 14, 576)	2,304	block_12_expand[...
block_12_expand_re... (ReLU)	(None, 14, 14, 576)	0	block_12_expand_...
block_12_depthwise (DepthwiseConv2D)	(None, 14, 14, 576)	5,184	block_12_expand_...
block_12_depthwise... (BatchNormalizatio...	(None, 14, 14, 576)	2,304	block_12_depthwi...
block_12_depthwise... (ReLU)	(None, 14, 14, 576)	0	block_12_depthwi...
block_12_project (Conv2D)	(None, 14, 14, 96)	55,296	block_12_depthwi...
block_12_project_BN (BatchNormalizatio...	(None, 14, 14, 96)	384	block_12_project...
block_12_add (Add)	(None, 14, 14, 96)	0	block_11_add[0][...] block_12_project...
block_13_expand (Conv2D)	(None, 14, 14, 576)	55,296	block_12_add[0][...
block_13_expand_BN (BatchNormalizatio...	(None, 14, 14, 576)	2,304	block_13_expand[...
block_13_expand_re... (ReLU)	(None, 14, 14, 576)	0	block_13_expand_...
block_13_pad (ZeroPadding2D)	(None, 15, 15, 576)	0	block_13_expand_...
block_13_depthwise (DepthwiseConv2D)	(None, 7, 7, 576)	5,184	block_13_pad[0][...
block_13_depthwise... (BatchNormalizatio...	(None, 7, 7, 576)	2,304	block_13_depthwi...
block_13_depthwise... (ReLU)	(None, 7, 7, 576)	0	block_13_depthwi...
block_13_project (Conv2D)	(None, 7, 7, 160)	92,160	block_13_depthwi...
block_13_project_BN (BatchNormalizatio...	(None, 7, 7, 160)	640	block_13_project...
block_14_expand (Conv2D)	(None, 7, 7, 960)	153,600	block_13_project...
block_14_expand_BN (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_14_expand[...
block_14_expand_re... (ReLU)	(None, 7, 7, 960)	0	block_14_expand_...

block_14_depthwise (DepthwiseConv2D)	(None, 7, 7, 960)	8,640	block_14_expand...
block_14_depthwise... (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_14_depthwi...
block_14_depthwise... (ReLU)	(None, 7, 7, 960)	0	block_14_depthwi...
block_14_project (Conv2D)	(None, 7, 7, 160)	153,600	block_14_depthwi...
block_14_project_BN (BatchNormalizatio...	(None, 7, 7, 160)	640	block_14_project...
block_14_add (Add)	(None, 7, 7, 160)	0	block_13_project... block_14_project...
block_15_expand (Conv2D)	(None, 7, 7, 960)	153,600	block_14_add[0][...
block_15_expand_BN (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_15_expand[...
block_15_expand_re... (ReLU)	(None, 7, 7, 960)	0	block_15_expand...
block_15_depthwise (DepthwiseConv2D)	(None, 7, 7, 960)	8,640	block_15_expand...
block_15_depthwise... (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_15_depthwi...
block_15_depthwise... (ReLU)	(None, 7, 7, 960)	0	block_15_depthwi...
block_15_project (Conv2D)	(None, 7, 7, 160)	153,600	block_15_depthwi...
block_15_project_BN (BatchNormalizatio...	(None, 7, 7, 160)	640	block_15_project...
block_15_add (Add)	(None, 7, 7, 160)	0	block_14_add[0][... block_15_project...
block_16_expand (Conv2D)	(None, 7, 7, 960)	153,600	block_15_add[0][...
block_16_expand_BN (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_16_expand[...
block_16_expand_re... (ReLU)	(None, 7, 7, 960)	0	block_16_expand...
block_16_depthwise (DepthwiseConv2D)	(None, 7, 7, 960)	8,640	block_16_expand...
block_16_depthwise... (BatchNormalizatio...	(None, 7, 7, 960)	3,840	block_16_depthwi...
block_16_depthwise... (ReLU)	(None, 7, 7, 960)	0	block_16_depthwi...
block_16_project (Conv2D)	(None, 7, 7, 320)	307,200	block_16_depthwi...
block_16_project_BN (BatchNormalizatio...	(None, 7, 7, 320)	1,280	block_16_project...
Conv_1 (Conv2D)	(None, 7, 7, 1280)	409,600	block_16_project...
Conv_1_bn (BatchNormalizatio...	(None, 7, 7, 1280)	5,120	Conv_1[0][0]
out_relu (ReLU)	(None, 7, 7, 1280)	0	Conv_1_bn[0][0]
global_average_poo... (GlobalAveragePool...	(None, 1280)	0	out_relu[0][0]
dense_2 (Dense)	(None, 256)	327,936	global_average_p...
dropout_1 (Dropout)	(None, 256)	0	dense_2[0][0]
dense_3 (Dense)	(None, 36)	9,252	dropout_1[0][0]

Total params: 2,595,172 (9.90 MB)

Trainable params: 337,188 (1.29 MB)

Non-trainable params: 2,257,984 (8.61 MB)

Epoch 1/15

126/126 ————— 155s 1s/step - accuracy: 0.0720 - loss: 3.5693 - val_accuracy: 0.1988 - val_loss: 3.0496

Epoch 2/15

126/126 ————— 189s 1s/step - accuracy: 0.1957 - loss: 2.9679 - val_accuracy: 0.3738 - val_loss: 2.5155

Epoch 3/15

126/126 ————— 156s 1s/step - accuracy: 0.2866 - loss: 2.5677 - val_accuracy: 0.5149 - val_loss: 2.0890

Epoch 4/15

```
# =====
# Part 16: Train ResNet50 Model
# باستخدام ResNet50 لتدريب نموذج
# =====

resnet_train_gen, resnet_val_gen = create_transfer_generators(resnet_preprocess)

resnet_base = ResNet50(
    weights="imagenet",
    include_top=False,
    input_shape=(TL_IMG_SIZE, TL_IMG_SIZE, 3)
)

resnet_model = build_transfer_learning_model(
    resnet_base,
    "ResNet50_ASL"
)


resnet_model.summary()

resnet_history = resnet_model.fit(
    resnet_train_gen,
    validation_data=resnet_val_gen,
    epochs=TL_EPOCHS
)

resnet_loss, resnet_acc = resnet_model.evaluate(
    resnet_val_gen,
    verbose=0
)

print(f"ResNet50 Validation Accuracy: {resnet_acc * 100:.2f}%")
print(f"ResNet50 Validation Loss : {resnet_loss:.4f}")

plot_training_curves(resnet_history, "ResNet50")
```



Found 2013 images belonging to 36 classes.

Found 503 images belonging to 36 classes.

Downloading data from https://storage.googleapis.com/tensorflow/keras-applications/resnet/resnet50_weights_tf_dim_ordering

94765736/94765736 3s 0us/step

Model: "ResNet50_ASL"

Layer (type)	Output Shape	Param #	Connected to
input_layer_2 (InputLayer)	(None, 224, 224, 3)	0	-
conv1_pad (ZeroPadding2D)	(None, 230, 230, 3)	0	input_layer_2[0]...
conv1_conv (Conv2D)	(None, 112, 112, 64)	9,472	conv1_pad[0][0]
conv1_bn (BatchNormalizatio...	(None, 112, 112, 64)	256	conv1_conv[0][0]
conv1_relu (Activation)	(None, 112, 112, 64)	0	conv1_bn[0][0]
pool1_pad (ZeroPadding2D)	(None, 114, 114, 64)	0	conv1_relu[0][0]
pool1_pool (MaxPooling2D)	(None, 56, 56, 64)	0	pool1_pad[0][0]
conv2_block1_1_conv (Conv2D)	(None, 56, 56, 64)	4,160	pool1_pool[0][0]
conv2_block1_1_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block1_1_c...
conv2_block1_1_relu (Activation)	(None, 56, 56, 64)	0	conv2_block1_1_b...
conv2_block1_2_conv (Conv2D)	(None, 56, 56, 64)	36,928	conv2_block1_1_r...
conv2_block1_2_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block1_2_c...
conv2_block1_2_relu (Activation)	(None, 56, 56, 64)	0	conv2_block1_2_b...
conv2_block1_0_conv (Conv2D)	(None, 56, 56, 256)	16,640	pool1_pool[0][0]
conv2_block1_3_conv (Conv2D)	(None, 56, 56, 256)	16,640	conv2_block1_2_r...
conv2_block1_0_bn (BatchNormalizatio...	(None, 56, 56, 256)	1,024	conv2_block1_0_c...
conv2_block1_3_bn (BatchNormalizatio...	(None, 56, 56, 256)	1,024	conv2_block1_3_c...
conv2_block1_add (Add)	(None, 56, 56, 256)	0	conv2_block1_0_b... conv2_block1_3_b...
conv2_block1_out (Activation)	(None, 56, 56, 256)	0	conv2_block1_add...
conv2_block2_1_conv (Conv2D)	(None, 56, 56, 64)	16,448	conv2_block1_out...
conv2_block2_1_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block2_1_c...
conv2_block2_1_relu (Activation)	(None, 56, 56, 64)	0	conv2_block2_1_b...
conv2_block2_2_conv (Conv2D)	(None, 56, 56, 64)	36,928	conv2_block2_1_r...
conv2_block2_2_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block2_2_c...
conv2_block2_2_relu (Activation)	(None, 56, 56, 64)	0	conv2_block2_2_b...
conv2_block2_3_conv (Conv2D)	(None, 56, 56, 256)	16,640	conv2_block2_2_r...
conv2_block2_3_bn (BatchNormalizatio...	(None, 56, 56, 256)	1,024	conv2_block2_3_c...
conv2_block2_add (Add)	(None, 56, 56, 256)	0	conv2_block1_out... conv2_block2_3_b...
conv2_block2_out (Activation)	(None, 56, 56, 256)	0	conv2_block2_add...
conv2_block3_1_conv (Conv2D)	(None, 56, 56, 64)	16,448	conv2_block2_out...

conv2_block3_1_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block3_1_c...
conv2_block3_1_relu (Activation)	(None, 56, 56, 64)	0	conv2_block3_1_b...
conv2_block3_2_conv (Conv2D)	(None, 56, 56, 64)	36,928	conv2_block3_1_r...
conv2_block3_2_bn (BatchNormalizatio...	(None, 56, 56, 64)	256	conv2_block3_2_c...
conv2_block3_2_relu (Activation)	(None, 56, 56, 64)	0	conv2_block3_2_b...
conv2_block3_3_conv (Conv2D)	(None, 56, 56, 256)	16,640	conv2_block3_2_r...
conv2_block3_3_bn (BatchNormalizatio...	(None, 56, 56, 256)	1,024	conv2_block3_3_c...
conv2_block3_add (Add)	(None, 56, 56, 256)	0	conv2_block2_out... conv2_block3_3_b...
conv2_block3_out (Activation)	(None, 56, 56, 256)	0	conv2_block3_add...
conv3_block1_1_conv (Conv2D)	(None, 28, 28, 128)	32,896	conv2_block3_out...
conv3_block1_1_bn (BatchNormalizatio...	(None, 28, 28, 128)	512	conv3_block1_1_c...
conv3_block1_1_relu (Activation)	(None, 28, 28, 128)	0	conv3_block1_1_b...
conv3_block1_2_conv (Conv2D)	(None, 28, 28, 128)	147,584	conv3_block1_1_r...
conv3_block1_2_bn (BatchNormalizatio...	(None, 28, 28, 128)	512	conv3_block1_2_c...
conv3_block1_2_relu (Activation)	(None, 28, 28, 128)	0	conv3_block1_2_b...
conv3_block1_0_conv (Conv2D)	(None, 28, 28, 512)	131,584	conv2_block3_out...
conv3_block1_3_conv (Conv2D)	(None, 28, 28, 512)	66,048	conv3_block1_2_r...
conv3_block1_0_bn (BatchNormalizatio...	(None, 28, 28, 512)	2,048	conv3_block1_0_c...
conv3_block1_3_bn (BatchNormalizatio...	(None, 28, 28, 512)	2,048	conv3_block1_3_c...
conv3_block1_add (Add)	(None, 28, 28, 512)	0	conv3_block1_0_b... conv3_block1_3_b...
conv3_block1_out (Activation)	(None, 28, 28, 512)	0	conv3_block1_add...
conv3_block2_1_conv (Conv2D)	(None, 28, 28, 128)	65,664	conv3_block1_out...
conv3_block2_1_bn (BatchNormalizatio...	(None, 28, 28, 128)	512	conv3_block2_1_c...
conv3_block2_1_relu (Activation)	(None, 28, 28, 128)	0	conv3_block2_1_b...
conv3_block2_2_conv (Conv2D)	(None, 28, 28, 128)	147,584	conv3_block2_1_r...
conv3_block2_2_bn (BatchNormalizatio...	(None, 28, 28, 128)	512	conv3_block2_2_c...
conv3_block2_2_relu (Activation)	(None, 28, 28, 128)	0	conv3_block2_2_b...
conv3_block2_3_conv (Conv2D)	(None, 28, 28, 512)	66,048	conv3_block2_2_r...
conv3_block2_3_bn (BatchNormalizatio...	(None, 28, 28, 512)	2,048	conv3_block2_3_c...
conv3_block2_add (Add)	(None, 28, 28, 512)	0	conv3_block1_out... conv3_block2_3_b...
conv3_block2_out (Activation)	(None, 28, 28, 512)	0	conv3_block2_add...
conv3_block3_1_conv (Conv2D)	(None, 28, 28, 128)	65,664	conv3_block2_out...
conv3_block3_1_bn (BatchNormalizatio...	(None, 28, 28, 128)	512	conv3_block3_1_c...

(BatchNormalization)	(None, 128)		
conv3_block3_1_relu (Activation)	(None, 28, 28, 128)	0	conv3_block3_1_b...
conv3_block3_2_conv (Conv2D)	(None, 28, 28, 128)	147,584	conv3_block3_1_r...
conv3_block3_2_bn (BatchNormalization)	(None, 28, 28, 128)	512	conv3_block3_2_c...
conv3_block3_2_relu (Activation)	(None, 28, 28, 128)	0	conv3_block3_2_b...
conv3_block3_3_conv (Conv2D)	(None, 28, 28, 512)	66,048	conv3_block3_2_r...
conv3_block3_3_bn (BatchNormalization)	(None, 28, 28, 512)	2,048	conv3_block3_3_c...
conv3_block3_add (Add)	(None, 28, 28, 512)	0	conv3_block2_out... conv3_block3_3_b...
conv3_block3_out (Activation)	(None, 28, 28, 512)	0	conv3_block3_add...
conv3_block4_1_conv (Conv2D)	(None, 28, 28, 128)	65,664	conv3_block3_out...
conv3_block4_1_bn (BatchNormalization)	(None, 28, 28, 128)	512	conv3_block4_1_c...
conv3_block4_1_relu (Activation)	(None, 28, 28, 128)	0	conv3_block4_1_b...
conv3_block4_2_conv (Conv2D)	(None, 28, 28, 128)	147,584	conv3_block4_1_r...
conv3_block4_2_bn (BatchNormalization)	(None, 28, 28, 128)	512	conv3_block4_2_c...
conv3_block4_2_relu (Activation)	(None, 28, 28, 128)	0	conv3_block4_2_b...
conv3_block4_3_conv (Conv2D)	(None, 28, 28, 512)	66,048	conv3_block4_2_r...
conv3_block4_3_bn (BatchNormalization)	(None, 28, 28, 512)	2,048	conv3_block4_3_c...
conv3_block4_add (Add)	(None, 28, 28, 512)	0	conv3_block3_out... conv3_block4_3_b...
conv3_block4_out (Activation)	(None, 28, 28, 512)	0	conv3_block4_add...
conv4_block1_1_conv (Conv2D)	(None, 14, 14, 256)	131,328	conv3_block4_out...
conv4_block1_1_bn (BatchNormalization)	(None, 14, 14, 256)	1,024	conv4_block1_1_c...
conv4_block1_1_relu (Activation)	(None, 14, 14, 256)	0	conv4_block1_1_b...
conv4_block1_2_conv (Conv2D)	(None, 14, 14, 256)	590,080	conv4_block1_1_r...
conv4_block1_2_bn (BatchNormalization)	(None, 14, 14, 256)	1,024	conv4_block1_2_c...
conv4_block1_2_relu (Activation)	(None, 14, 14, 256)	0	conv4_block1_2_b...
conv4_block1_0_conv (Conv2D)	(None, 14, 14, 1024)	525,312	conv3_block4_out...
conv4_block1_3_conv (Conv2D)	(None, 14, 14, 1024)	263,168	conv4_block1_2_r...
conv4_block1_0_bn (BatchNormalization)	(None, 14, 14, 1024)	4,096	conv4_block1_0_c...
conv4_block1_3_bn (BatchNormalization)	(None, 14, 14, 1024)	4,096	conv4_block1_3_c...
conv4_block1_add (Add)	(None, 14, 14, 1024)	0	conv4_block1_0_b... conv4_block1_3_b...
conv4_block1_out (Activation)	(None, 14, 14, 1024)	0	conv4_block1_add...
conv4_block2_1_conv (Conv2D)	(None, 14, 14, 256)	262,400	conv4_block1_out...
conv4_block2_1_bn (BatchNormalization)	(None, 14, 14, 256)	1,024	conv4_block2_1_c...
conv4_block2_1_relu (Activation)	(None, 14, 14, 256)	0	conv4_block2_1_b...

conv4_block2_1_conv (Activation)	(None, 14, 14, 256)	590,080	conv4_block2_1_r...
conv4_block2_2_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block2_2_c...
conv4_block2_2_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block2_2_b...
conv4_block2_2_relu (Activation)	(None, 14, 14, 256)	263,168	conv4_block2_2_r...
conv4_block2_3_conv (Conv2D)	(None, 14, 14, 1024)	4,096	conv4_block2_3_c...
conv4_block2_3_bn (BatchNormalizatio...	(None, 14, 14, 1024)	0	conv4_block1_out...
conv4_block2_add (Add)	(None, 14, 14, 1024)	0	conv4_block2_3_b...
conv4_block2_out (Activation)	(None, 14, 14, 1024)	262,400	conv4_block2_add...
conv4_block3_1_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block2_out...
conv4_block3_1_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block3_1_c...
conv4_block3_1_relu (Activation)	(None, 14, 14, 256)	590,080	conv4_block3_1_b...
conv4_block3_2_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block3_1_r...
conv4_block3_2_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block3_2_c...
conv4_block3_2_relu (Activation)	(None, 14, 14, 256)	263,168	conv4_block3_2_b...
conv4_block3_3_conv (Conv2D)	(None, 14, 14, 1024)	4,096	conv4_block3_2_r...
conv4_block3_3_bn (BatchNormalizatio...	(None, 14, 14, 1024)	0	conv4_block3_3_c...
conv4_block3_add (Add)	(None, 14, 14, 1024)	0	conv4_block2_out...
conv4_block3_out (Activation)	(None, 14, 14, 1024)	0	conv4_block3_3_b...
conv4_block3_add...			
conv4_block4_1_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block3_out...
conv4_block4_1_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block4_1_c...
conv4_block4_1_relu (Activation)	(None, 14, 14, 256)	590,080	conv4_block4_1_b...
conv4_block4_2_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block4_1_r...
conv4_block4_2_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block4_2_c...
conv4_block4_2_relu (Activation)	(None, 14, 14, 256)	263,168	conv4_block4_2_b...
conv4_block4_3_conv (Conv2D)	(None, 14, 14, 1024)	4,096	conv4_block4_2_r...
conv4_block4_3_bn (BatchNormalizatio...	(None, 14, 14, 1024)	0	conv4_block4_3_c...
conv4_block4_add (Add)	(None, 14, 14, 1024)	0	conv4_block3_out...
conv4_block4_out (Activation)	(None, 14, 14, 1024)	0	conv4_block4_3_b...
conv4_block4_add...			
conv4_block5_1_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block4_out...
conv4_block5_1_bn (BatchNormalizatio...	(None, 14, 14, 256)	0	conv4_block5_1_c...
conv4_block5_1_relu (Activation)	(None, 14, 14, 256)	590,080	conv4_block5_1_b...
conv4_block5_2_conv (Conv2D)	(None, 14, 14, 256)	1,024	conv4_block5_1_r...
conv4_block5_2_bn (BatchNormalizatio...	(None, 14, 14, 256)		conv4_block5_2_c...

conv4_block5_2_relu (Activation)	(None, 14, 14, 256)	0	conv4_block5_2_b...
conv4_block5_3_conv (Conv2D)	(None, 14, 14, 1024)	263,168	conv4_block5_2_r...
conv4_block5_3_bn (BatchNormalizatio...	(None, 14, 14, 1024)	4,096	conv4_block5_3_c...
conv4_block5_add (Add)	(None, 14, 14, 1024)	0	conv4_block4_out... conv4_block5_3_b...
conv4_block5_out (Activation)	(None, 14, 14, 1024)	0	conv4_block5_add...
conv4_block6_1_conv (Conv2D)	(None, 14, 14, 256)	262,400	conv4_block5_out...
conv4_block6_1_bn (BatchNormalizatio...	(None, 14, 14, 256)	1,024	conv4_block6_1_c...
conv4_block6_1_relu (Activation)	(None, 14, 14, 256)	0	conv4_block6_1_b...
conv4_block6_2_conv (Conv2D)	(None, 14, 14, 256)	590,080	conv4_block6_1_r...
conv4_block6_2_bn (BatchNormalizatio...	(None, 14, 14, 256)	1,024	conv4_block6_2_c...
conv4_block6_2_relu (Activation)	(None, 14, 14, 256)	0	conv4_block6_2_b...
conv4_block6_3_conv (Conv2D)	(None, 14, 14, 1024)	263,168	conv4_block6_2_r...
conv4_block6_3_bn (BatchNormalizatio...	(None, 14, 14, 1024)	4,096	conv4_block6_3_c...
conv4_block6_add (Add)	(None, 14, 14, 1024)	0	conv4_block5_out... conv4_block6_3_b...
conv4_block6_out (Activation)	(None, 14, 14, 1024)	0	conv4_block6_add...
conv5_block1_1_conv (Conv2D)	(None, 7, 7, 512)	524,800	conv4_block6_out...
conv5_block1_1_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block1_1_c...
conv5_block1_1_relu (Activation)	(None, 7, 7, 512)	0	conv5_block1_1_b...
conv5_block1_2_conv (Conv2D)	(None, 7, 7, 512)	2,359,808	conv5_block1_1_r...
conv5_block1_2_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block1_2_c...
conv5_block1_2_relu (Activation)	(None, 7, 7, 512)	0	conv5_block1_2_b...
conv5_block1_0_conv (Conv2D)	(None, 7, 7, 2048)	2,099,200	conv4_block6_out...
conv5_block1_3_conv (Conv2D)	(None, 7, 7, 2048)	1,050,624	conv5_block1_2_r...
conv5_block1_0_bn (BatchNormalizatio...	(None, 7, 7, 2048)	8,192	conv5_block1_0_c...
conv5_block1_3_bn (BatchNormalizatio...	(None, 7, 7, 2048)	8,192	conv5_block1_3_c...
conv5_block1_add (Add)	(None, 7, 7, 2048)	0	conv5_block1_0_b... conv5_block1_3_b...
conv5_block1_out (Activation)	(None, 7, 7, 2048)	0	conv5_block1_add...
conv5_block2_1_conv (Conv2D)	(None, 7, 7, 512)	1,049,088	conv5_block1_out...
conv5_block2_1_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block2_1_c...
conv5_block2_1_relu (Activation)	(None, 7, 7, 512)	0	conv5_block2_1_b...
conv5_block2_2_conv (Conv2D)	(None, 7, 7, 512)	2,359,808	conv5_block2_1_r...
conv5_block2_2_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block2_2_c...
conv5_block2_2_relu (Activation)	(None, 7, 7, 512)	0	conv5_block2_2_b...

conv5_block2_3_conv (Conv2D)	(None, 7, 7, 2048)	1,050,624	conv5_block2_2_r...
conv5_block2_3_bn (BatchNormalizatio...	(None, 7, 7, 2048)	8,192	conv5_block2_3_c...
conv5_block2_add (Add)	(None, 7, 7, 2048)	0	conv5_block1_out...
conv5_block2_out (Activation)	(None, 7, 7, 2048)	0	conv5_block2_add...
conv5_block3_1_conv (Conv2D)	(None, 7, 7, 512)	1,049,088	conv5_block2_out...
conv5_block3_1_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block3_1_c...
conv5_block3_1_relu (Activation)	(None, 7, 7, 512)	0	conv5_block3_1_b...
conv5_block3_2_conv (Conv2D)	(None, 7, 7, 512)	2,359,808	conv5_block3_1_r...
conv5_block3_2_bn (BatchNormalizatio...	(None, 7, 7, 512)	2,048	conv5_block3_2_c...
conv5_block3_2_relu (Activation)	(None, 7, 7, 512)	0	conv5_block3_2_b...
conv5_block3_3_conv (Conv2D)	(None, 7, 7, 2048)	1,050,624	conv5_block3_2_r...
conv5_block3_3_bn (BatchNormalizatio...	(None, 7, 7, 2048)	8,192	conv5_block3_3_c...
conv5_block3_add (Add)	(None, 7, 7, 2048)	0	conv5_block2_out...
conv5_block3_out (Activation)	(None, 7, 7, 2048)	0	conv5_block3_add...
global_average_poo... (GlobalAveragePool...	(None, 2048)	0	conv5_block3_out...
dense_4 (Dense)	(None, 256)	524,544	global_average_p...
dropout_2 (Dropout)	(None, 256)	0	dense_4[0][0]
dense_5 (Dense)	(None, 36)	9,252	dropout_2[0][0]

Total params: 24,121,508 (92.02 MB)

Trainable params: 533,796 (2.04 MB)

Non-trainable params: 23,587,712 (89.98 MB)

Epoch 1/15

126/126 ————— 506s 4s/step - accuracy: 0.0691 - loss: 3.5842 - val_accuracy: 0.1968 - val_loss: 3.0298

Epoch 2/15

126/126 ————— 491s 4s/step - accuracy: 0.1893 - loss: 2.9203 - val_accuracy: 0.3817 - val_loss: 2.5572

Epoch 3/15

126/126 ————— 495s 4s/step - accuracy: 0.2896 - loss: 2.5228 - val_accuracy: 0.5129 - val_loss: 2.1782

Epoch 4/15

126/126 ————— 495s 4s/step - accuracy: 0.3616 - loss: 2.2541 - val_accuracy: 0.5586 - val_loss: 1.9032

Epoch 5/15

126/126 ————— 496s 4s/step - accuracy: 0.4158 - loss: 2.0181 - val_accuracy: 0.5865 - val_loss: 1.6963

Epoch 6/15

126/126 ————— 496s 4s/step - accuracy: 0.4471 - loss: 1.8536 - val_accuracy: 0.6322 - val_loss: 1.5231

Epoch 7/15

126/126 ————— 495s 4s/step - accuracy: 0.5042 - loss: 1.7067 - val_accuracy: 0.7078 - val_loss: 1.3232

Epoch 8/15

126/126 ————— 543s 4s/step - accuracy: 0.5460 - loss: 1.5437 - val_accuracy: 0.7256 - val_loss: 1.2548

Epoch 9/15

126/126 ————— 499s 4s/step - accuracy: 0.5683 - loss: 1.4643 - val_accuracy: 0.7614 - val_loss: 1.1340

Epoch 10/15

126/126 ————— 499s 4s/step - accuracy: 0.5991 - loss: 1.3574 - val_accuracy: 0.7356 - val_loss: 1.0917

Epoch 11/15

126/126 ————— 499s 4s/step - accuracy: 0.6195 - loss: 1.2891 - val_accuracy: 0.7614 - val_loss: 0.9766

Epoch 12/15

126/126 ————— 502s 4s/step - accuracy: 0.6428 - loss: 1.2191 - val_accuracy: 0.7913 - val_loss: 0.9543

Epoch 13/15

126/126 ————— 501s 4s/step - accuracy: 0.6617 - loss: 1.1341 - val_accuracy: 0.8032 - val_loss: 0.8565

Epoch 14/15


126/126 ————— 502s 4s/step - accuracy: 0.6761 - loss: 1.0769 - val_accuracy: 0.7952 - val_loss: 0.8252

Epoch 15/15

126/126 ————— 497s 4s/step - accuracy: 0.6975 - loss: 1.0089 - val_accuracy: 0.8151 - val_loss: 0.7910

ResNet50 Validation Accuracy: 81.51%

ResNet50 Validation Loss : 0.7910



```
# =====
# Part 17: Train EfficientNetB0 Model
# باستخدام EfficientNetB0 لتدريب نموذج
# =====

efficientnet_train_gen, efficientnet_val_gen = create_transfer_generators(efficientnet_preprocess)

efficientnet_base = EfficientNetB0(
    weights="imagenet",
    include_top=False,
    input_shape=(TL_IMG_SIZE, TL_IMG_SIZE, 3)
)

efficientnet_model = build_transfer_learning_model(
    efficientnet_base,
    "EfficientNetB0_AS_L"
)

efficientnet_model.summary()

efficientnet_history = efficientnet_model.fit(
    efficientnet_train_gen,
    validation_data=efficientnet_val_gen,
    epochs=TL_EPOCHS
)

efficientnet_loss, efficientnet_acc = efficientnet_model.evaluate(
    efficientnet_val_gen,
    verbose=0
)

print(f"EfficientNetB0 Validation Accuracy: {efficientnet_acc * 100:.2f}%")
print(f"EfficientNetB0 Validation Loss : {efficientnet_loss:.4f}")

plot_training_curves(efficientnet_history, "EfficientNetB0")
```

```
# =====
# Part 18: Create Models Comparison Table
# إنشاء جدول مقارنة بين CNN و MobileNetV2 و ResNet50 و EfficientNetB0
# =====

cnn_best_val_acc = max(cnn_history.history["val_accuracy"])
cnn_best_val_loss = min(cnn_history.history["val_loss"])

mobilenet_best_val_acc = max(mobilenet_history.history["val_accuracy"])
mobilenet_best_val_loss = min(mobilenet_history.history["val_loss"])

resnet_best_val_acc = max(resnet_history.history["val_accuracy"])
resnet_best_val_loss = min(resnet_history.history["val_loss"])

efficientnet_best_val_acc = max(efficientnet_history.history["val_accuracy"])
efficientnet_best_val_loss = min(efficientnet_history.history["val_loss"])

comparison_table = pd.DataFrame({
    "Model": [
        "Custom CNN",
        "MobileNetV2",
        "ResNet50",
        "EfficientNetB0"
    ],
    "Best Validation Accuracy": [
        f"{cnn_best_val_acc * 100:.2f}%",
        f"{mobilenet_best_val_acc * 100:.2f}%",
        f"{resnet_best_val_acc * 100:.2f}%",
        f"{efficientnet_best_val_acc * 100:.2f}%"
    ]
})
```