PROJECT ABSTRACT

Food Calorie Estimation from Photos using CNN

A two-headed CNN estimates dish identity and calories from meal photos: an EfficientNet-B0 backbone with a 101-class dish head (Food-101) and a calorie regression head. Delivered built-to-order with training notebook, nutrition table, weights, evaluation report, web demo, PPT and viva kit.

01. Title

Food Calorie Estimation from Photos using CNN.

02. Introduction / Background

Diet-logging apps fail when users must type every meal manually; photo-based estimation removes that friction. The task splits naturally into two coupled problems: recognizing the dish, and estimating its energy content. The Food-101 dataset (Bossard, Guillaumin and Van Gool, ETH Zurich) provides 101,000 labeled food images across 101 dishes — the standard classification benchmark — while calorie values come from curated per-dish nutrition tables. A shared convolutional backbone with two heads (classification + regression) learns visual features useful to both tasks at once.

03. Problem Statement

Manual calorie tracking has high dropout because every meal must be searched and portion-estimated by hand. A photo-based estimator must handle real-world messiness: mixed dishes, occlusions, unknown portions and hidden oil. This project builds an honest prototype: dish classification on Food-101 plus a calorie regression head trained on portion-annotated samples, with a user-selected portion multiplier and a clearly stated error band — a logging aid, explicitly not medical advice.

04. Objectives

- Train a dish classifier on Food-101 (101 classes) using an EfficientNet-B0 backbone.
- Add a calorie regression head predicting kcal per 100 g from shared backbone features.
- Build a per-dish nutrition table (protein/carbs/fat per 100 g) joined by class label.
- Evaluate dish top-1 accuracy and calorie MAE on the held-out test split.
- Deliver a web demo: upload a meal photo, pick a portion, get dish + calorie/macro estimate.

05. Proposed Methodology

(1) The EfficientNet-B0 backbone is initialized from ImageNet and fine-tuned on Food-101's 75,750 training images with augmentation. (2) Two heads are trained jointly: Dense(101)+softmax with cross-entropy for the dish, and Dense(1)+ReLU with MSE for kcal/100 g, combined as a weighted loss (0.6·CE + 0.4·MSE). (3) Portion-annotated samples teach the regression head scale; at inference the user selects a portion multiplier (0.5×/1×/2×). (4) Macros are computed from the nutrition table for the predicted dish. Evaluation reports top-1 accuracy and calorie MAE on the 25,250-image test split.

06. System Architecture / Working

A meal photo is resized to 224×224 and normalized. The EfficientNet backbone extracts multi-scale food features (texture, color, shape). The classification head outputs the dish with confidence; the regression head outputs kcal per 100 g. Multiplying by the dish's reference weight and the user's portion factor gives the serving estimate, and the nutrition table supplies protein/carbs/fat. The demo displays the dish confidence alongside the calorie number so users see when the estimate rests on an uncertain classification — the key honesty feature of the design.

07. Hardware / Software Requirements

- Python 3.10+, TensorFlow/Keras or PyTorch
- Food-101 dataset (~5 GB) + curated nutrition table (ships with project)
- GPU recommended; 8 GB RAM
- Flask demo backend; HTML/CSS/JS frontend

08. Expected Outcomes

A two-headed food model, measured dish top-1 accuracy and calorie MAE on the held-out test split, the nutrition table, and a working demo. The report states the measured error band plainly and documents failure modes (occluded portions, mixed thalis) — exactly the material a viva examiner probes.

09. Applications

- Diet-logging and fitness apps
- Canteen/campus nutrition awareness kiosks
- Dietician client onboarding (rough intake estimates)

10. Future Scope

- Automatic portion-size estimation from depth/reference objects
- Multi-dish segmentation for mixed plates (thali)
- Personalized targets with user history tracking