

# Investigating the Impact of Large-Scale Pre-training on Nutritional Content Estimation from 2D Images

Michele Andrade<sup>1</sup>, Guilherme A. L. Silva<sup>2</sup>, Valéria Santos<sup>1</sup>,  
Gladston Moreira<sup>1</sup> and Eduardo Luz<sup>1</sup>

<sup>1</sup>Computing Department, Federal University of Ouro Preto, Ouro Preto, Brazil, 35402-136

<sup>2</sup>Postgraduate Program in Computer Science, Federal University of Ouro Preto, Ouro Preto, Brazil, 35402-136

**Abstract**—Estimating the nutritional content of food from images is a critical task with significant implications for health and dietary monitoring. This is challenging, especially when relying solely on 2D images, due to the variability in food presentation, lighting, and the inherent difficulty in inferring volume and mass without depth information. Furthermore, reproducibility in this domain is hampered by the reliance of state-of-the-art methods on proprietary datasets for large-scale pre-training. In this paper, we investigate the impact of large-scale pre-training datasets on the performance of deep learning models for nutritional estimation using only 2D images. We fine-tune and evaluate Vision Transformer (ViT) models pre-trained on two large public datasets, ImageNet and COYO, comparing their performance against baseline CNN models (InceptionV2 and ResNet-50) and a state-of-the-art method pre-trained on the proprietary JFT-300M dataset. We conduct extensive experiments on the Nutrition5k dataset, a large-scale collection of real-world food plates with high-precision nutritional annotations. Our evaluation using Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAE%) reveals that models pre-trained on JFT-300M significantly outperform those pre-trained on public datasets. Unexpectedly, the model pre-trained on the massive COYO dataset performs worse than the model pre-trained on ImageNet for this specific regression task, refuting our initial hypothesis. Our analysis provides quantitative evidence highlighting the critical role of pre-training dataset characteristics, including scale, domain relevance, and curation quality, for effective transfer learning in 2D nutritional estimation.

**Index Terms**—Nutritional Estimation, Food Recognition, Deep Learning, Vision Transformer, Pre-training, Transfer Learning, Computer Vision.

## I. INTRODUCTION

Estimating the nutritional content of food solely from images presents a significant challenge at the intersection of computer vision and nutrition science. The growing global focus on health, dietary management, and chronic disease prevention has amplified the need for accessible and accurate tools to track nutrient intake. The ability to determine calories, carbohydrates, proteins, and fats from a simple photograph of a meal would revolutionize personal health monitoring, dietary studies, and nutritional support systems.

This task is inherently complex despite its potential, particularly when relying on standard 2D image data from ubiquitous devices like smartphone cameras. Key difficulties include the immense variability in food presentation (lighting, angle, plating), the challenge of distinguishing individual ingredients within mixed dishes, and critically, the ambiguity in estimating serving size and volume from a single 2D

perspective. Traditional methods often require manual input or supplementary data (e.g., food scales, measuring tools), limiting scalability and user convenience.

Recent research has attempted to tackle these challenges through diverse approaches. Datasets like Nutrition5k [1], pic2kcal [2], ISIA Food-500 [3], and Food2K [4] have emerged, providing large collections of food images, often with associated nutritional information and ingredient lists. Some studies have explored the use of depth sensors [1], [5] or explicit volume estimation techniques [6] to mitigate the limitations of 2D data. Deep learning models, particularly Convolutional Neural Networks (CNNs), have shown promise in recognizing food items and even predicting nutritional values directly from images [2], [7]–[9]. However, many state-of-the-art approaches, such as the baseline presented with the Nutrition5k dataset [1], achieve their top performance by leveraging models pre-trained on massive, often proprietary, datasets like JFT-300M [10].

This reliance on inaccessible pre-training data poses a significant challenge to the reproducibility of research in the community. It restricts widespread adoption, further development, and fair comparison of new methods. Motivated by this, our work focuses on investigating the effectiveness of publicly available large-scale pre-training datasets for nutritional estimation from 2D images, specifically exploring the potential of the Vision Transformer (ViT) architecture [11] due to its strong performance on large-scale vision tasks.

This study aims to answer the following research questions:

- Can Vision Transformer models effectively estimate nutritional content (calories, mass, macronutrients) from 2D food images when fine-tuned on a specific food dataset?
- How does pre-training a ViT model on a massive, publicly available multimodal dataset (COYO) compare to pre-training on a widely used visual dataset (ImageNet) for this regression task?
- Can pre-training on public datasets (ImageNet or COYO) achieve performance comparable to models pre-trained on large proprietary datasets (JFT-300M)?

Based on the scale of the COYO dataset, which is closer in size to JFT-300M than ImageNet, we hypothesize that pre-training on COYO would yield superior results compared to pre-training on ImageNet and potentially bridge the performance gap with the JFT-300M baseline for 2D nutritional estimation.

We formulate the problem as a multi-task regression task predicting total calories, mass, fat, carbohydrates, and pro-

teins. We implement and evaluate ViT models pre-trained on ImageNet and COYO, fine-tuning them on the Nutrition5k dataset [1]. We compare their performance using MAE and MAE% against standard CNN baselines (InceptionV2, ResNet-50) trained under similar conditions and the reported results of the state-of-the-art baseline from [1].

Our main contributions are:

- Adaptation and evaluation of the Vision Transformer architecture for nutritional content regression from 2D food images;
- Empirical investigation into the impact of pre-training on large-scale public datasets (ImageNet and COYO) for this specific regression problem, comparing them directly;
- Quantitative comparison of ViT performance against established CNN architectures and a state-of-the-art baseline pre-trained on a proprietary dataset, highlighting the challenges of achieving competitive results without access to such data;
- Providing insights into which characteristics of large pre-training datasets (scale, domain, curation) appear most critical for successful transfer learning in 2D nutritional estimation.

## II. RELATED WORK

Developing robust food recognition and nutritional estimation methods relies heavily on the availability of large, diverse, and accurately annotated datasets. Several efforts have focused on collecting food image datasets, often with varying levels of annotation detail. Datasets such as pic2kcal [2] and ISIA Food-500 [3] were constructed by scraping images and associated information from websites, providing large-scale collections. Pic2kcal includes recipes and nutritional details, while ISIA Food-500 focuses on fine-grained food category recognition across 500 categories. Food2K [4] further expands this by collecting over 1 million images for 2000 categories, demonstrating broad applicability in tasks like recognition and retrieval. While these datasets are valuable for food category recognition, their reliance on crowdsourced or web-scraped nutritional data can introduce inaccuracies, and they often lack precise information about serving sizes or ingredient weights for specific dishes depicted in images.

Addressing the crucial challenge of portion size estimation, the Nutrition5k dataset [1] provides a unique resource by collecting 5,000 real-world food plates, along with accompanying videos, depth images, and precise component weights, thereby enabling highly accurate nutritional annotations. This structured data collection approach contrasts with web-scraping and provides a more reliable ground truth for nutritional estimation. Similarly, Naritomi et al. [5] explored the use of 3D scanning to create 3D models of food items and plates, aiding in volume estimation and highlighting the importance of 3D information. Our work utilizes the Nutrition5k dataset, focusing specifically on its 2D components to investigate models applicable in scenarios lacking depth sensors.

Various deep learning architectures have been applied to food recognition and nutritional tasks. CNNs, such as VGG,

Inception, and ResNet, have been widely used for food image classification [8] and as backbones for nutritional prediction models [1], [2], [9]. Multi-task learning approaches, where a single model simultaneously predicts multiple attributes like ingredients, categories, and nutritional values, have shown benefits by exploiting correlations between tasks [2], [9]. For instance, Ruede et al. [2] used a multi-task setup with CNNs on pic2kcal to predict calories and macronutrients. Thames et al. [1] employed a multi-task regression head on an InceptionV2 base for calorie, mass, and macronutrient prediction on Nutrition5k.

More recently, Vision Transformers (ViT) [11], inspired by NLP success, have demonstrated state-of-the-art performance in image recognition when pre-trained on massive datasets. Min et al. [4] leveraged transformer-like attention mechanisms in their “deep progressive region enhancement network” for food recognition and retrieval on Food2K. While transformers show promise in vision, their application to nutritional regression from images, and the impact of different large-scale pre-training sources on this specific task, remains less explored. Critically, the superior performance of the Nutrition5k baseline [1] is attributed to pre-training on the proprietary JFT-300M dataset, which is inaccessible for most researchers.

Our work directly compares established CNNs (InceptionV2, ResNet-50) with the Vision Transformer for 2D nutritional estimation on Nutrition5k. Unlike methods relying on 3D data, we focus on the challenging 2D scenario. Our primary contribution lies in empirically evaluating the effectiveness of public large-scale pre-training datasets (ImageNet and COYO) for ViT on this regression task, contrasting their performance and characteristics against the JFT-300M pre-trained baseline to shed light on the reproducibility gap and the factors influencing successful transfer learning in this domain.

## III. BACKGROUND

The problem of estimating nutritional content from images is approached using supervised learning techniques, specifically deep neural networks, which are highly effective at learning complex patterns from visual data [12]. Our task is formulated as a regression problem, aiming to predict continuous numerical values (mass, calories, and macronutrients) directly from a 2D image.

### A. Problem Setting: Nutritional Estimation as Regression

We formalize the problem of nutritional content estimation from a single 2D image as a multi-task regression problem. Given an input food image  $I \in \mathbb{R}^{H \times W \times C}$ , where  $H$ ,  $W$ , and  $C$  are the height, width, and number of channels (e.g., 3 for RGB), the objective is to predict a vector  $\hat{Y} = (\hat{y}_w, \hat{Y}_m, \hat{y}_{cal})$ , where:

- $\hat{y}_w \in \mathbb{R}$  is the estimated total mass (in grams) of the plate;
- $\hat{Y}_m \in \mathbb{R}^k$  is a vector of estimated masses (in grams) for  $k$  specific macronutrients. In this work,  $k = 3$  for fat, carbohydrates, and protein, so  $\hat{Y}_m = (\hat{y}_{fat}, \hat{y}_{carb}, \hat{y}_{protein})$ ;
- $\hat{y}_{cal} \in \mathbb{R}$  is the estimated total caloric content (in kcal) of the plate.

The ground truth for a given image  $I_i$  is  $Y_i = (y_i^w, Y_i^m, y_i^{cal})$ , derived from measured weights of individual ingredients and their known nutritional values per gram. The task is to learn a mapping function  $f : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{k+2}$  that minimizes a defined loss function between the predicted values  $\hat{Y}_i = f(I_i)$  and the ground truth  $Y_i$  over a dataset  $D = \{I_i, Y_i\}_{i=1}^N$ .

### B. Deep Learning Models

We investigate three distinct deep learning architectures commonly used in computer vision and adapted for regression: InceptionV2, ResNet-50, and the Vision Transformer (ViT).

1) *InceptionV2*: InceptionV2 [13] improves upon the original Inception architecture by factorizing large convolutions into smaller or asymmetric ones to enhance computational efficiency while maintaining representation power. It employs parallel branches with different filter sizes within inception modules.

2) *ResNet-50*: ResNet-50 [14] is a 50-layer network utilizing residual connections (“skip connections”) to alleviate the degradation problem in deep networks. These connections facilitate gradient flow, enabling the effective training of models that are significantly deeper than standard feedforward architectures.

3) *Vision Transformer (ViT)*: The Vision Transformer (ViT) [11] adapts the Transformer architecture [15] for image processing. It treats an image as a sequence of patches, processes them using self-attention mechanisms in a Transformer encoder, and lacks the strong inductive biases of CNNs, relying heavily on large-scale pre-training data to achieve high performance. As illustrated in fig. 1, input images are divided into patches, embedded, and processed by a standard Transformer encoder.

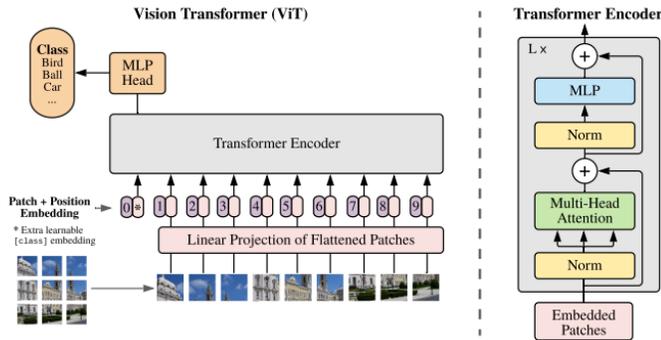


Fig. 1: Vision Transformer (ViT) Architecture. (Source: [11])

### C. Multi-task Learning Head

Following the approach in [1], we employ a multi-task learning head on top of the backbone CNN (InceptionV2, ResNet-50) or ViT model. This head shares initial fully connected layers across tasks before splitting into task-specific final layers. For CNNs, the output feature map from an intermediate layer (specifically, ‘mixed5c’ for InceptionV2 and the layer before global average pooling for ResNet-50) is spatially pooled (3x3 average pooling with stride 2). For ViT,

the output embedding of the classification token from the last layer is used directly.

A sequence of two shared fully connected (FC) layers, each with 4096 dimensions, processes the pooled features and token embedding. Subsequently, separate task-specific FC layers (also with 4096 dimensions) are used for mass, calories, and a combined head for macronutrients. Each task has a final FC layer mapping to its respective output dimension (1 for mass, calories, 3 for macronutrients).

The overall loss function is a summation of the losses for each task, trained jointly:

$$\mathcal{L}(I, Y|W) = \mathcal{L}_m(I, Y^m|W) + \mathcal{L}_c(I, y^{cal}|W) + \mathcal{L}_w(I, y^w|W) \quad (1)$$

where  $W$  represents the model parameters,  $I$  is the input image, and  $Y = (y^w, Y^m, y^{cal})$  are the ground truth labels. For the CNN experiments compared to the baseline, we use Mean Absolute Error (MAE) as the loss for each task, consistent with [1]:

$$\begin{aligned} \mathcal{L}_m(I, Y^m|W) &= \frac{1}{|M|} \sum_{j \in M} |\hat{y}_j^m(I|W) - y_j^m| \\ \mathcal{L}_c(I, y^{cal}|W) &= |\hat{y}^{cal}(I|W) - y^{cal}| \\ \mathcal{L}_w(I, y^w|W) &= |\hat{y}^w(I|W) - y^w| \end{aligned} \quad (2)$$

where  $M = \{\text{fat, carbohydrate, protein}\}$ ,  $\hat{y}(I|W)$  denotes the model’s prediction for input  $I$  with parameters  $W$ , and  $y$  denotes the ground truth label. For ViT experiments, we used Mean Squared Error (MSE) as the loss function for each task, a common choice for regression problems that was found to work well empirically during initial testing. The overall loss function for ViT training thus becomes:

$$\begin{aligned} \mathcal{L}_{ViT}(I, Y|W) &= \text{MSE}(\hat{Y}_m(I|W), Y_m) + \text{MSE}(\hat{y}_{cal}(I|W), y_{cal}) \\ &\quad + \text{MSE}(\hat{y}_w(I|W), y_w) \end{aligned} \quad (3)$$

where  $\text{MSE}(a, b) = \frac{1}{N} \sum_{i=1}^N (a_i - b_i)^2$  for vectors  $a, b$  of size  $N$ , or  $(a - b)^2$  for scalars.

Figure 2 illustrates the overall approach of estimating nutritional content from a 2D image using a deep learning model with a multi-task approach.

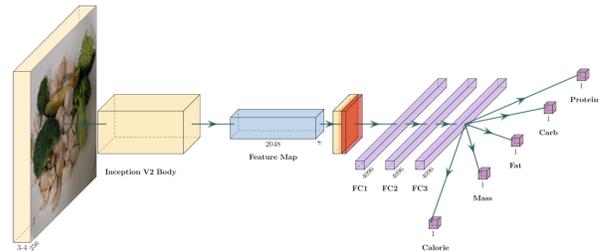


Fig. 2: Overall approach for nutritional content estimation from a 2D food image using a deep learning model. (Adapted from [1])

## IV. EXPERIMENTAL SETUP

### A. Dataset

Our experiments are conducted using the Nutrition5k dataset [1]. This dataset was specifically designed for nutritional understanding from images, providing high-precision annotations. It comprises approximately 5,000 unique real-world food plates, built from over 250 distinct ingredients. The data collection process involved a custom sensor system that weighed and scanned each plate component individually (fig. 3). Each plate is annotated with total weight, weights for individual components, and derived nutritional values (calories, fat, carbs, protein) calculated from USDA data. The dataset also includes RGB videos and aerial RGB-D images for many plates.

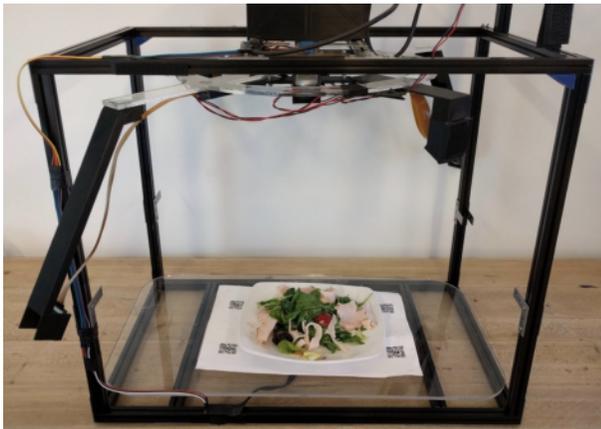


Fig. 3: Data collection equipment used for the Nutrition5k dataset. (Source: [1])

For this study, which focuses on the challenging task of nutritional estimation from 2D images only, we utilize the RGB video data provided in Nutrition5k. We extract static frames from these videos to simulate standard 2D image input. Specifically, we sample one frame every five frames from the recorded videos associated with each plate.

The dataset is split following the original paper’s protocol: 90% of the plates are used for training, and the remaining 10% are held out for testing. This split is performed at the plate level to ensure that no images from the same plate appear in both the training and testing sets, thereby providing a realistic evaluation of generalization. The dataset exhibits variability in food types, portions, and quantities, reflecting real-world culinary diversity, as illustrated by the most common ingredients by mass (table I).

### B. Models and Training Details

We implement and train three types of models: InceptionV2, ResNet-50, and Vision Transformer (ViT).

**CNN Baselines (InceptionV2 and ResNet-50):** We implemented InceptionV2 based on the description in [1] and ResNet-50 [14] using standard architectures. Both models were initialized with weights pre-trained on the ImageNet dataset [16]. Input images were resized and center-cropped to  $256 \times 256$  pixels. Training utilized the RMSprop optimizer with a learning

TABLE I: Frequency of ingredient usage (in thousands)

Ingredient	Value (k)
Egg whites	47.8
Chicken	47.7
Scrambled eggs	40.3
Olives	37.3
Cauliflower	29.6
Broccoli	29.3
Lemon	28.0
Berries	24.3
Fish	21.3
Apple	19.4
Cheese pizza	19.2
Sweet potato	18.3
Carrot	18.2
Pineapple	17.6
Caesar salad	17.2
Cantaloupe	16.3
Cherry tomatoes	15.7
Beef	15.5
Pork	15.3
Salmon	14.8
Sausage	14.0
Roasted potatoes	13.7
Asparagus	13.6
Brussels sprouts	13.3
White rice	13.2
Mixed greens	12.9
Squash	11.6
Corn on the cob	10.4
Cucumbers	10.2
Watermelon	10.2

rate of  $1 \times 10^{-4}$ , momentum 0.9, decay 0.9, and epsilon 1.0. The multi-task head described in section III-C was attached, and models were trained using the L1 loss function for each task.

**Vision Transformer (ViT):** We used the CLIP ViT architecture available from the Hugging Face Transformers library. The pre-trained classification head was removed, and a linear regression head was added on top of the output embedding of the classification token from the last Transformer layer. This head maps to 5 output values: total mass, total fat, total carbs, total protein, and total calories.

We explored two different pre-training sources for the ViT model:

- ViT pre-trained on ImageNet: A ViT model pre-trained on the standard ImageNet-21k dataset.
- ViT pre-trained on COYO: A ViT model pre-trained on the large-scale COYO-700M dataset [17], which contains over 700 million image-text pairs.

For fine-tuning on Nutrition5k, both ViT models were trained using the Adam optimizer with an empirically tuned learning rate. The multi-task MSE loss function was used. Training was conducted for two different durations to observe convergence behavior: 60 epochs and 300 epochs. Batch Normalization was applied in the regression head.

### C. Evaluation Metrics

To evaluate the performance of the models in estimating total mass, calories, and individual macronutrient masses (fat,

carbs, protein), we use the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAE%).

The Mean Absolute Error (MAE) is defined as:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (4)$$

where  $n$  is the number of samples in the test set,  $y_i$  is the ground truth value for the  $i$ -th sample, and  $\hat{y}_i$  is the predicted value for the  $i$ -th sample.

The Mean Absolute Percentage Error (MAE%) provides a relative measure of error:

$$MAE\% = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100\% \quad (5)$$

Note that MAE% can be unstable or undefined for ground truth values close to zero. However, for the aggregated metrics (total calories, total mass, total macronutrient masses), ground truth values are generally non-zero and substantial.

Lower values for both MAE and MAE% indicate better model performance.

## V. RESULTS AND DISCUSSION

This section presents the experimental results and discusses their implications, particularly addressing our research questions and hypothesis regarding the impact of pre-training datasets on 2D nutritional estimation.

### A. CNN Baselines vs. Nutrition5k Baseline

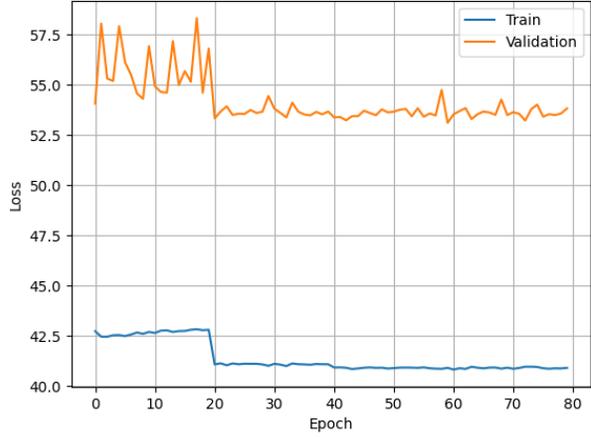
Table II presents the performance of our implemented InceptionV2 and ResNet-50 models, both pre-trained on ImageNet and fine-tuned on Nutrition5k using L1 loss, compared to the reported performance of the Nutrition5k baseline [1], which used an InceptionV2-based architecture pre-trained on JFT-300M.

TABLE II: Comparison of MAE and MAE% for InceptionV2 and ResNet-50 (ImageNet pre-trained) against the Nutrition5k Baseline (JFT-300M pre-trained).

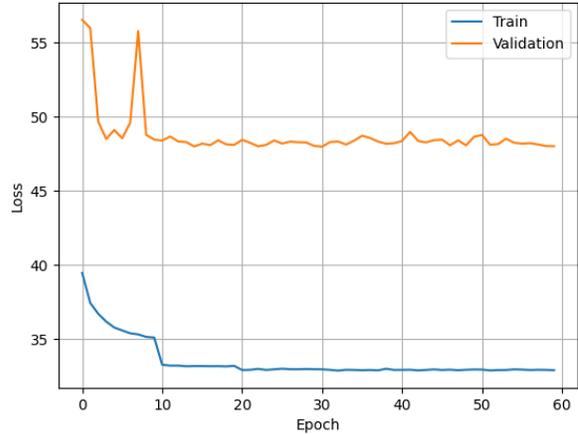
Metric	InceptionV2 (ImageNet)		ResNet-50 (ImageNet)		Baseline (JFT-300M)	
	MAE	MAE %	MAE	MAE %	MAE	MAE %
Total Calories	140.88	55.90	117.38	46.57	<b>70.6</b>	<b>26.1</b>
Total Mass	82.50	43.62	66.83	35.33	<b>40.4</b>	<b>18.8</b>
Total Fat	9.93	73.00	9.02	66.32	<b>5.0</b>	<b>34.2</b>
Total Carb	11.19	58.48	9.80	51.21	<b>6.1</b>	<b>31.9</b>
Total Protein	10.54	67.14	8.92	56.80	<b>5.5</b>	<b>29.5</b>

As observed, the ResNet-50 model consistently outperforms the InceptionV2 model across all metrics when both are pre-trained on the ImageNet dataset. This suggests that the ResNet architecture’s ability to train deeper models effectively through residual connections provides an advantage for this task, even compared to the InceptionV2 variant used in the baseline. Figure 4 shows the training and validation loss curves, indicating relatively stable convergence for both models, although ResNet-50’s performance is better overall.

However, despite ResNet-50’s better performance compared to InceptionV2, both models fall significantly short of the Nutrition5k baseline in terms of MAE and MAE%. The



(a) InceptionV2 (L1 loss)



(b) ResNet-50 (L1 loss)

Fig. 4: Training and validation loss curves for InceptionV2 and ResNet-50 models trained with L1 loss.

baseline achieves substantially lower errors across all nutritional parameters. This stark difference highlights the impact of the pre-training dataset. The baseline was pre-trained on JFT-300M [10], a proprietary dataset with 300 million images and 375 million labels. This massive scale and diversity, encompassing a broad range of visual concepts, provide a much richer starting point for fine-tuning compared to ImageNet (approximately 14 million images). This initial comparison underscores the challenge of replicating state-of-the-art results in vision tasks when the original models leverage inaccessible pre-training data.

### B. Vision Transformer (ViT) Experiments

The results of the ViT experiments are presented, specifically aimed at evaluating the impact of different large-scale public pre-training datasets (ImageNet and COYO) on 2D nutritional

estimation, and assessing their effectiveness in approaching the performance of the JFT-300M baseline.

Table III and Table IV show the performance of ViT models pre-trained on COYO and ImageNet, trained for 60 and 300 epochs, respectively, compared to the Nutrition5k baseline.

TABLE III: Comparison of MAE and MAE% for ViT models (60 Epochs) against the Nutrition5k Baseline.

Metric	ViT (COYO)		ViT (ImageNet)		Baseline (JFT-300M)	
	MAE	MAE %	MAE	MAE %	MAE	MAE %
Total Calories	144.40	57.29	<b>87.99</b>	<b>34.91</b>	<b>70.6</b>	<b>26.1</b>
Total Mass	91.98	48.63	<b>63.87</b>	<b>33.77</b>	<b>40.4</b>	<b>18.8</b>
Total Fat	10.48	77.04	<b>7.22</b>	<b>53.07</b>	<b>5.0</b>	<b>34.2</b>
Total Carb	12.19	63.68	<b>8.28</b>	<b>43.28</b>	<b>6.1</b>	<b>31.9</b>
Total Protein	11.47	73.03	<b>6.63</b>	<b>42.24</b>	<b>5.5</b>	<b>29.5</b>

TABLE IV: Comparison of MAE and MAE% for ViT models (300 Epochs) against the Nutrition5k Baseline.

Metric	ViT (COYO)		ViT (ImageNet)		Baseline (JFT-300M)	
	MAE	MAE %	MAE	MAE %	MAE	MAE %
Total Calories	136.97	54.34	<b>95.29</b>	<b>37.81</b>	<b>70.6</b>	<b>26.1</b>
Total Mass	82.23	43.48	<b>60.55</b>	<b>32.01</b>	<b>40.4</b>	<b>18.8</b>
Total Fat	10.38	76.27	<b>7.40</b>	<b>54.43</b>	<b>5.0</b>	<b>34.2</b>
Total Carb	11.42	59.64	<b>8.29</b>	<b>43.30</b>	<b>6.1</b>	<b>31.9</b>
Total Protein	10.21	65.05	<b>6.86</b>	<b>43.67</b>	<b>5.5</b>	<b>29.5</b>

Consistent with the CNN results, the Nutrition5k baseline pre-trained on JFT-300M remains superior to both ViT models pre-trained on public datasets across all metrics and training durations. This further reinforces the significant advantage conferred by pre-training on extremely large, diverse, and presumably well-curated datasets like JFT-300M.

Crucially, the results for the ViT models pre-trained on COYO and ImageNet directly address our second and third research questions and test our hypothesis. The ViT model pre-trained on ImageNet significantly outperforms the ViT model pre-trained on COYO in nearly all metrics, for both 60 and 300 epochs. For example, the MAE for Total Calories after 300 epochs is 136.97 for COYO ViT vs 95.29 for ImageNet ViT. This refutes our hypothesis that pre-training on COYO would be superior to ImageNet due to its larger size and purported similarity to JFT-300M. The ImageNet pre-trained ViT consistently yields lower errors, demonstrating better transfer learning capability for this task compared to the COYO pre-trained model.

Several factors could explain the poorer performance of the COYO pre-trained ViT compared to ImageNet. While COYO is massive ( 700M image-text pairs), it is a multimodal dataset collected from the web, potentially containing significant noise and less structured annotations compared to ImageNet, which is specifically curated for object classification. More importantly, the domain of images in COYO might be less relevant to the specific visual characteristics of food plates compared to the broad range of objects and scenes in ImageNet, which includes many food categories. Although JFT-300M is also web-scale, Sun et al. [10] found that its 300M images and 375M labels, despite some noise, provide tremendous visual diversity, leading to representations that generalize exceptionally well.

Our results suggest that for the task of nutritional estimation from food images, the curated diversity and focus of datasets like ImageNet (or the extreme scale and breadth of JFT-300M) are more beneficial than the sheer size and multimodal nature of COYO.

Increasing the training duration from 60 to 300 epochs generally improved the performance of both ViT models, with MAE values decreasing for most metrics (e.g., Calories MAE for COYO: 144.40 at 60 epochs vs 136.97 at 300 epochs; for ImageNet: 87.99 at 60 epochs vs 95.29 at 300 epochs - Note: ImageNet calories MAE slightly increased at 300 epochs, while mass/macro MAE decreased. This oscillation is visible in loss curves). Figure 5 illustrates the training and validation loss curves for the 300-epoch runs. The ImageNet ViT exhibits a lower training loss and somewhat more stable validation loss compared to the COYO ViT, indicating better learning on the training data and slightly improved generalization, although both models show some volatility in validation loss, which is typical when fine-tuning large models on smaller datasets. While increased training time helps, it does not bridge the gap to the JFT-300M baseline performance.

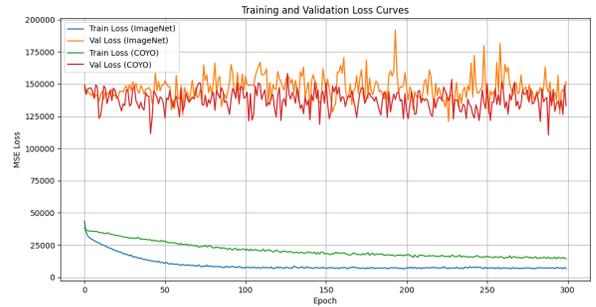


Fig. 5: Training and validation loss curves for ViT models pre-trained on COYO and ImageNet (300 Epochs).

### C. Discussion on Research Questions

Regarding our first research question, ViT models, when fine-tuned, can indeed estimate nutritional content from 2D food images, achieving performance significantly better than random but still considerably lower than the JFT-300M baseline. Their effectiveness is heavily dependent on the pre-training data. For the second question, pre-training on ImageNet resulted in demonstrably better performance for this regression task compared to pre-training on COYO, contrary to our hypothesis. And addressing the third question, models pre-trained on public datasets (ImageNet and COYO), even the massive COYO, were unable to match the performance of the baseline pre-trained on the proprietary JFT-300M dataset. This highlights the significant advantage that large-scale, possibly more curated and domain-relevant (in terms of visual diversity) proprietary data provides, posing a major hurdle for reproducibility in this field. The results quantitatively demonstrate that achieving state-of-the-art performance requires access to data resources that are not universally available.

The source code and experimental results are available at <https://github.com/michele-andrade/Nutricao-Inteligente>.

## VI. CONCLUSION

In this work, we investigated the problem of estimating nutritional content from 2D food images, focusing on the impact of large-scale pre-training datasets on model performance.

We implemented and evaluated Vision Transformer models pre-trained on the public ImageNet and COYO datasets, comparing them against CNN baselines (InceptionV2, ResNet-50) and a state-of-the-art baseline pre-trained on the proprietary JFT-300M dataset, using the Nutrition5k dataset.

Our experiments yielded clear insights. Firstly, while ViT models can be effectively fine-tuned for nutritional regression, their performance is strongly tied to the pre-training data. Secondly, contrary to our initial hypothesis, pre-training on the massive COYO dataset did not surpass pre-training on ImageNet for this specific task; the ImageNet pre-trained ViT consistently performed better. Finally, and most significantly, neither of the models pre-trained on public datasets could match the performance of the JFT-300M pre-trained baseline.

These findings quantitatively demonstrate that the scale and, more importantly, the specific characteristics (domain relevance, curation quality) of pre-training data are paramount for achieving state-of-the-art results in 2D nutritional estimation. The performance gap observed with the JFT-300M baseline underscores the reproducibility challenge faced by researchers without access to such proprietary resources. Our results suggest that simply using any large public dataset may not be sufficient; the nature of the visual concepts learned during pre-training matters critically for transfer learning to this complex regression task.

## ACKNOWLEDGMENT

This work was supported by the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq, grants 308400/2022-4, 307151/2022-0), Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES - grant 001), Fundação de Amparo à Pesquisa do Estado de Minas Gerais (FAPEMIG, grant APQ-01647-22). We also thank the Universidade Federal de Ouro Preto (UFOP) for their support.

## REFERENCES

- [1] Q. Thames, A. Karpur, W. Norris, F. Xia, L. Panait, T. Weyand, and J. Sim, "Nutrition5k: Towards automatic nutritional understanding of generic food," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8903–8911.
- [2] R. Ruede, V. Heusser, L. Frank, A. Roitberg, M. Haurilet, and R. Stiefelhagen, "Multi-task learning for calorie prediction on a novel large-scale recipe dataset enriched with nutritional information," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 4001–4008.
- [3] W. Min, L. Liu, Z. Wang, Z. Luo, X. Wei, X. Wei, and S. Jiang, "Isia food-500: A dataset for large-scale food recognition via stacked global-local attention network," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 393–401.
- [4] W. Min, Z. Wang, Y. Liu, M. Luo, L. Kang, X. Wei, X. Wei, and S. Jiang, "Large scale visual food recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [5] S. Naritomi and K. Yanai, "Hungry networks: 3d mesh reconstruction of a dish and a plate from a single dish image for estimating food volume," in *Proceedings of the 2nd ACM International Conference on Multimedia in Asia*, 2021, pp. 1–7.
- [6] A. Meyers, N. Johnston, V. Rathod, A. Korattikara, A. Gorban, N. Silberman, S. Guadarrama, G. Papandreou, J. Huang, and K. P. Murphy, "Im2calories: towards an automated mobile vision food diary," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1233–1241.
- [7] J. Chen and C.-W. Ngo, "Deep-based ingredient recognition for cooking recipe retrieval," in *Proceedings of the 24th ACM international conference on Multimedia*, 2016, pp. 32–41.
- [8] R. Yunus, O. Arif, H. Afzal, M. F. Amjad, H. Abbas, H. N. Bokhari, S. T. Haider, N. Zafar, and R. Nawaz, "A framework to estimate the nutritional value of food in real time using deep learning techniques," *IEEE Access*, vol. 7, pp. 2643–2652, 2018.
- [9] J. Chen, B. Zhu, C.-W. Ngo, T.-S. Chua, and Y.-G. Jiang, "A study of multi-task and region-wise deep learning for food ingredient recognition," *IEEE Transactions on Image Processing*, vol. 30, pp. 1514–1526, 2020.
- [10] C. Sun, A. Shrivastava, S. Singh, and A. Gupta, "Revisiting unreasonable effectiveness of data in deep learning era," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 843–852.
- [11] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [12] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [13] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [16] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [17] M. Byeon, B. Park, H. Kim, S. Lee, W. Baek, and S. Kim, "Coyo-700m: Image-text pair dataset," <https://github.com/kakaobrain/coyo-dataset>, 2022.