

PortionNet: Distilling 3D Geometric Knowledge for Food Nutrition Estimation

Darrin Bright¹, Rakshith Raj¹, Kanchan Keisham¹

¹ Vellore Institute of Technology, India

{darrin.bright2022, rakshith.raj2022}@vitstudent.ac.in, kanchan.keisham@vit.ac.in

Abstract

Accurate food nutrition estimation from single images is challenging due to the loss of 3D information. While depth-based methods provide reliable geometry, they remain inaccessible on most smartphones because of depth-sensor requirements. To overcome this challenge, we propose PortionNet, a novel cross-modal knowledge distillation framework that learns geometric features from point clouds during training while requiring only Red, Green, Blue (RGB) images at inference. Our approach employs a dual-mode training strategy where a lightweight adapter network mimics point cloud representations, enabling pseudo-3D reasoning without any specialized hardware requirements. PortionNet achieves state-of-the-art performance on MetaFood3D, outperforming all previous methods in both volume and energy estimation. Cross-dataset evaluation on SimpleFood45 further demonstrates strong generalization in energy estimation.

1. Introduction

Dietary monitoring plays an important role in managing chronic diseases like obesity and diabetes, yet manual food logging suffers from high error rates [3]. Image-based methods offer a solution to replace manual logging, but accurately estimating nutrition from a single RGB image is difficult due to the absence of geometric features.

Depth sensor methods address this issue by providing reliable 3D geometry [4], but these methods are limited by the cost and accessibility of hardware requirements like LiDAR on smartphones. Therefore, RGB-only methods that do not require specialized hardware can be utilized on smartphones. However, these methods have difficulties with spatial reasoning, as current state-of-the-art RGB-only methods exhibit a mean absolute percentage error (MAPE) of 68.05% [14]. This significant accuracy gap has prevented RGB-only methods from practical deployment.

Recent advances in knowledge distillation [7, 8] enable cross-modal capability transfer. However, existing food es-

timination methods either require depth sensors during inference or fail to utilize 3D effectively. No prior work has explored distilling 3D spatial reasoning into RGB-only networks for food nutrition estimation.

To address this gap, we propose PortionNet, a novel cross-modal knowledge distillation framework that learns 3D-aware representations from point clouds during training but requires only RGB during inference. A Point Cloud Geometry Encoder serves as a teacher, while a lightweight RGB-to-Geometry Adapter learns to mimic its feature space. Dual-mode training alternates between multi-modal (RGB + point cloud) and RGB-only batches, limiting over-reliance on 3D information while encouraging strong pseudo-3D feature learning.

Our contributions can be summarized as follows:

- We propose PortionNet, a novel cross-modal knowledge distillation framework that allows RGB models to learn 3D geometric features from point clouds during training, eliminating depth sensor requirements at inference.
- We introduce a dual-mode training strategy with a lightweight RGB-to-Geometry Adapter that learns to generate pseudo-3D features for accurate geometric reasoning from standard RGB images.
- Our method achieves state-of-the-art performance on MetaFood3D [1] with 17.43% volume MAPE and 15.36% energy MAPE. Cross-dataset evaluation on SimpleFood45 [20] achieves state-of-the-art energy estimation (12.17% MAPE), demonstrating effective generalization.

2. Related Work

Early food portion estimation systems relied heavily on geometric primitives and manual calibration [5, 23], limiting their applicability to real-world scenarios. RGB-D methods [4, 10, 11, 13] substantially improved accuracy by utilizing explicit depth information for 3D reconstruction. However, these methods require specialized depth sensors such as LiDAR or structured light cameras, which remain unavailable on most consumer smartphones.

Modern RGB-only methods attempt to infer 3D information from 2D images through various strategies.

Transformer-based methods [6, 19] use self-attention mechanisms for fine-grained food recognition, while depth prediction approaches [16] synthesize pseudo-depth maps to guide volume estimation. Multi-scale networks [22] combine segmentation and regression for nutrient prediction. Lately, 3D model-based methods [20] use existing 3D food models for portion estimation from single images. Despite these advances, RGB-only methods face significant challenges in recovering 3D geometry from monocular images, where scale information is fundamentally lost. MFP3D [14], the current state-of-the-art RGB-only method, achieves 41.43% volume MAPE and 68.05% energy MAPE.

Knowledge distillation [8] enables student models to mimic stronger teachers by matching their learned representations. Cross-modal distillation [7] extends this pattern to transfer knowledge between different input modalities, such as from depth to RGB. Recent work explores bidirectional learning [9], attention-based feature alignment [21], and handling incomplete modalities [12, 24], where certain sensors are unavailable at inference. However, these methods typically focus on tasks where semantic understanding is enough. Food nutrition estimation uniquely requires both semantic recognition (food type and composition) and metric geometric reasoning (absolute volume). To our knowledge, no prior work has explored distilling 3D spatial reasoning from point clouds into RGB-only networks for this task.

3. Methodology

3.1. Architecture Overview

PortionNet learns 3D geometric representations by using point-cloud features during training while requiring only RGB images at inference. As shown in Figure 1, the framework consists of: (1) dual RGB encoders that extract visual features, (2) an augmented PointNet geometry encoder that serves as the teacher, (3) an RGB-to-point-cloud adapter that maps RGB features to geometry-aware space, and (4) a fusion module with task-specific prediction heads.

3.2. Multi-Modal Feature Extraction

3.2.1. Dual RGB Encoders

We extract visual features using pretrained ViT-B/16 and ResNet-18 encoders, concatenating their outputs (1280-d) before projecting to 256 dimensions via a two-layer MLP, resulting in feature vectors of shape $[B \times 256]$.

3.2.2. PointNet Geometry Encoder

We use PointNet [15] to extract geometric features from point clouds, which learns spatial structure from unordered 3D data. We make two modifications: (1) append learned embeddings of bounding box dimensions ($\mathbb{R}^3 \rightarrow \mathbb{R}^{64}$) for

explicit scale information, and (2) apply multi-scale adaptive pooling at 64, 128, 256, 512 resolutions before global aggregation. The final representation is projected to 256 dimensions.

3.2.3. RGB-to-Point-Cloud Adapter

The adapter is a lightweight two-layer MLP that maps the extracted RGB features into a geometry-aware embedding space. This adapter enables RGB-only inference by mimicking PointNet’s output, eliminating the need for sensors while preserving geometric reasoning.

3.3. Cross-Modal Knowledge Distillation

3.3.1. Dual-Mode Training

Training alternates between multimodal mode using both RGB and PointNet features (70%) and RGB-only mode using RGB and adapter-generated features (30%), preventing over-reliance on real point clouds while ensuring strong pseudo-3D learning.

3.3.2. Distillation Objective

We supervise the adapter to mimic PointNet’s geometric features. Let $f_{\text{adapter}} \in \mathbb{R}^{256}$ and $f_{\text{PointNet}} \in \mathbb{R}^{256}$ denote the feature embeddings produced by the RGB adapter and the PointNet encoder, respectively. The distillation loss combines three complementary objectives:

$$\mathcal{L}_{\text{distill}} = w_{\text{mse}} \mathcal{L}_{\text{MSE}} + w_{\text{cos}} \mathcal{L}_{\text{cos}} + w_{\text{KL}} \mathcal{L}_{\text{KL}} \quad (1)$$

where $w_{\text{mse}} = 0.7$, $w_{\text{cos}} = 0.2$, $w_{\text{KL}} = 0.1$. The MSE loss

$$\mathcal{L}_{\text{MSE}} = \|f_{\text{adapter}} - f_{\text{PointNet}}\|_2^2 \quad (2)$$

matches feature magnitudes, the cosine loss

$$\mathcal{L}_{\text{cos}} = 1 - \text{COS}(f_{\text{adapter}}, f_{\text{PointNet}}) \quad (3)$$

aligns feature directions, and the KL divergence loss

$$\mathcal{L}_{\text{KL}} = \text{KL}(\sigma(f_{\text{PointNet}}/T) \parallel \sigma(f_{\text{adapter}}/T)) \quad (4)$$

matches probability distributions with softmax σ and temperature $T = 4$.

3.4. Feature Fusion and Prediction

3.4.1. Cross-Modal Fusion

RGB and point cloud features are fused using bidirectional cross-attention with residual connections and layer normalization. The attended features are concatenated and processed through a fusion MLP.

3.4.2. Task-Specific Prediction Heads

We employ three prediction heads: (1) 3-layer MLP for the 108-class food classification, (2) volume estimation with Softplus activation, ensuring positive outputs, and (3) energy estimation using concatenated features to model physical volume-energy dependency.

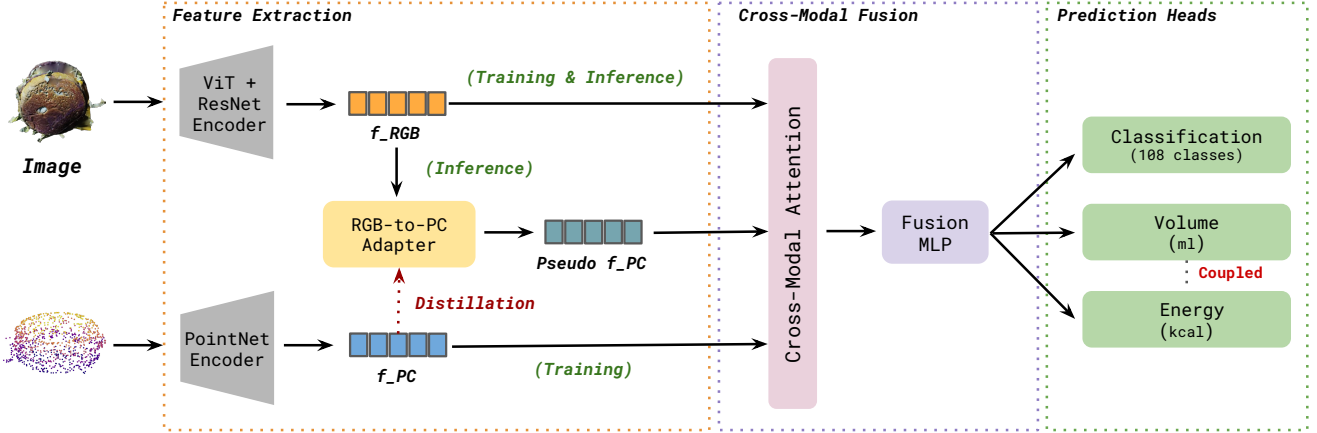


Figure 1. Overview of the proposed PortionNet framework.

3.5. Training Objective

The overall training objective combines three weighted loss components:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{cls}}\mathcal{L}_{\text{cls}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{distill}}\mathcal{L}_{\text{distill}} \quad (5)$$

where \mathcal{L}_{cls} is the classification loss, \mathcal{L}_{reg} is the regression loss for volume and energy prediction, $\mathcal{L}_{\text{distill}}$ is the feature distillation loss (Eq. 1), and λ_{cls} , λ_{reg} , λ_{distill} are their respective task weights.

Classification Loss. We employ cross-entropy with label smoothing ($\epsilon = 0.05$):

$$\mathcal{L}_{\text{cls}} = - \sum_{i=1}^C y_i \log(\hat{y}_i) \quad (6)$$

where $C = 108$, y_i is the smoothed label, and \hat{y}_i is predicted probability.

Regression Loss. We employ a combination of L1 and Huber losses ($\delta = 0.5$) for both volume and energy prediction. Let v and \hat{v} denote the ground-truth and predicted food volume, and e and \hat{e} denote the ground-truth and predicted energy values. The regression loss is calculated as:

$$\mathcal{L}_{\text{reg}} = 0.4(\text{L1}(v, \hat{v}) + \text{Huber}(v, \hat{v}, \delta = 0.5)) + 0.6(\text{L1}(e, \hat{e}) + \text{Huber}(e, \hat{e}, \delta = 0.5)) \quad (7)$$

where \hat{e} is batch-normalized for scale invariance.

3.6. Implementation Details

We train PortionNet using the AdamW optimizer with differential learning rates (1×10^{-4} for pretrained encoders, 5×10^{-4} for heads and adapter), OneCycleLR scheduling (10% warmup, 25 epochs), an effective batch size of 64 (16×4 accumulation), 8-head cross-modal attention, and gradient clipping (norm 1.0). We employ GradNorm [2] with $\lambda_{\text{reg}} = 0.1$ and $\lambda_{\text{distill}} = 0.5$ in multimodal mode.

Table 1. MetaFood3D Results: RGB-only vs RGB+PC modes

Metric	RGB	RGB+PC
R^2	0.926 ± 0.004	0.926 ± 0.005
Accuracy (%)	98.34 ± 0.58	98.26 ± 0.60
Vol. MAE (mL)	25.39 ± 2.59	25.16 ± 2.41
Vol. MAPE (%)	17.43 ± 0.81	17.33 ± 0.84
Eng. MAE (kcal)	32.26 ± 0.43	32.64 ± 0.66
Eng. MAPE (%)	15.36 ± 1.33	15.39 ± 1.46

4. Results

We evaluate PortionNet on MetaFood3D [1], containing 637 objects across 108 classes with 12,752 RGB images and point clouds. We follow MFP3D’s 80/20 split. For cross-dataset evaluation, we use SimpleFood45 [20] with 513 images across 12 classes. All results are averaged over 3 seeds with standard deviations.

4.1. Evaluation on MetaFood3D

Table 1 shows that PortionNet achieves 98.34% classification accuracy and strong regression (17.43% volume MAPE and 15.36% energy MAPE) using only standard RGB cameras for inference. The minimal gap between RGB-only and RGB+PC modes validates effective knowledge distillation.

4.2. Comparison with RGB-based Methods

Table 2 shows that PortionNet achieves substantial improvements over the state-of-the-art RGB-only method MFP3D [14], reducing volume MAPE from 41.43% to 17.43% (57.9% relative improvement) and energy MAPE from 68.05% to 15.36% (77.4% relative improvement), demonstrating effective knowledge distillation without requiring depth sensors at inference.

Table 2. Comparison with other methods on MetaFood3D

Method	Vol. MAPE (%) ↓	Eng. MAPE (%) ↓
Density Map [17] [†]	–	663.43
Stereo [4] [†]	210.90	–
Voxel [18] [†]	104.07	–
3D Assisted [20] [†]	79.33	102.25
MFP3D [14] [†]	41.43	68.05
PortionNet (Ours)	17.43 ± 0.81 (-24.0)	15.36 ± 1.33 (-52.7)

[†]Results from MFP3D [14]

Table 3. Cross-dataset generalization on SimpleFood45

Method	Vol. MAPE (%) ↓	Eng. MAPE (%) ↓
Density Map Only [†]	–	159.48
Density Map Summing [†]	–	93.16
Voxel Reconstruction [†]	24.51	–
3D Assisted [†]	14.01	25.13
MFP3D [†]	16.15	24.03
PortionNet (Ours)	23.51 ± 0.92 (+9.50)	12.17 ± 1.36 (-11.86)

[†]Results from MFP3D [14]

Table 4. Baseline and adapter ablation on MetaFood3D

Configuration	Vol. MAE (mL) ↓	Eng. MAE (kcal) ↓
Baseline (no adapter/distill)*	29.18	36.37
PortionNet (RGB)	21.74	29.03
PortionNet (RGB+PC)	21.22	29.13
No Adapter (direct RGB) [†]	27.92	38.66
With Adapter (PortionNet)	26.97	33.70

*Dual encoders + PointNet with concatenation.

[†]Trained without distillation, inference bypasses the adapter.

4.3. Cross-Dataset Generalization

We also perform cross-data generalization on the SimpleFood45 dataset as shown in Table 3. PortionNet achieves 12.17% energy MAPE, a 49.4% relative improvement over MFP3D [14] (24.03%), demonstrating strong cross-dataset generalization.

Volume estimation yields a 23.51% MAPE, compared to MFP3D’s 16.15%. This asymmetry occurs because energy depends on learned food composition features that transfer across datasets, while volume requires absolute spatial measurements tied to the camera parameters of a dataset.

4.4. Ablation Studies

We conduct ablation studies on MetaFood3D using a single seed with validation-optimized hyperparameters.

4.4.1. Baseline and Adapter Comparison

We use a baseline without the adapter or distillation. Table 4 shows PortionNet achieves 25.5% volume and 20.2%

Table 5. Hyperparameter ablation on MetaFood3D

Configuration	Vol. MAE (mL) ↓	Eng. MAE (kcal) ↓
<i>RGB-only Training Ratio (α), $\lambda=0.5$</i>		
$\alpha=0.0$	27.57	34.82
$\alpha=0.3$	21.74	29.03
$\alpha=0.5$	27.13	39.69
<i>Distillation Weight (λ), $\alpha=0.3$</i>		
$\lambda=0.0$	37.07	37.00
$\lambda=0.5$	21.74	29.03
$\lambda=1.0$	23.62	36.32

RGB-only inference.

energy improvement. Removing distillation ($\lambda_{\text{distill}} = 0$) causes significant degradation: 33.2% and 28.4% increases in energy and volume MAE, confirming that the adapter requires explicit distillation supervision for effective learning.

4.4.2. Hyperparameter Optimization

We validate two critical hyperparameters (Table 5). The **RGB-only training ratio** (α) determines the proportion of batches using adapter-generated versus real point cloud features. Insufficient ($\alpha = 0$) or excessive ($\alpha = 0.5$) ratios degrade performance, while optimal $\alpha = 0.3$ achieves 21% error reduction. The **distillation weight** (λ) guides the feature alignment strength between the adapter and PointNet. Ablating distillation ($\lambda = 0$) causes 70% and 28% degradation in volume and energy, while optimal $\lambda = 0.5$ demonstrates effective transfer with 41% and 22% improvements. Excessive weighting ($\lambda = 1.0$) over-prioritizes feature matching, degrading the performance.

5. Conclusion

We presented PortionNet, a novel cross-modal knowledge distillation framework that enables accurate nutrition estimation from RGB images without depth sensors. By training a lightweight adapter to mimic point cloud geometric features, our approach achieves significant improvements over prior methods while requiring only a single image at inference.

Our results show that cross-modal distillation can effectively transfer geometric knowledge from 3D to 2D representations, enabling nutrition estimation on standard devices without specialized hardware.

For future work, one promising direction is to address the volume estimation domain shift across imaging devices by developing camera-agnostic geometric representations through computational photography techniques which could provide consistent scale estimation independent of camera parameters.

References

- [1] Yuhao Chen, Jiangpeng He, Chris Czarnecki, Gautham Vinod, Talha Ibn Mahmud, Siddeshwar Raghavan, Jinge Ma, Dayou Mao, Saejith Nair, Pengcheng Xi, Alexander Wong, Edward Delp, and Fengqing Zhu. Metafood3d: Large 3d food object dataset with nutrition values. *arXiv preprint arXiv:2409.01966*, 2024. 1, 3
- [2] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *ICML*, 2018. 3
- [3] Phawinpon Chotwanvirat, Aree Prachansuwan, Pimnapant Sridonpai, and Wantanee Kriengsinyos. Advancements in using ai for dietary assessment based on food images: Scoping review. *J. Med. Internet Res.*, 26:e51432, 2024. 1
- [4] Joachim Dehais, Marios Anthimopoulos, Sergey Shevchik, and Stavroula Mougiakakou. Two-view 3d reconstruction for food volume estimation. *IEEE Trans. Multimedia*, 19(5): 1090–1099, 2017. 1, 4
- [5] Shaobo Fang, Chang Liu, Fengqing Zhu, Edward J. Delp, and Carol J. Boushey. Single-view food portion estimation based on geometric models. In *IEEE Int. Symp. Multimedia*, pages 385–390. IEEE, 2015. 1
- [6] Wei Feng, Jie Li, Hui Wang, and Xiaoming Zhang. Fine-grained food image recognition using laplacian pyramids and bilinear pooling. In *Int. Conf. Multimedia Model.*, pages 256–268. Springer, 2023. 2
- [7] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *CVPR*, pages 2827–2836, 2016. 1, 2
- [8] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2
- [9] Fushuo Huo, Wenchao Zhang, Jianbing Mao, and Miaoqing Wang. C2kd: Bridging the modality gap for cross-modal knowledge distillation. In *CVPR*, 2024. 2
- [10] Ji-hwan Kim, Dong-seok Lee, and Soon-kak Kwon. Food classification and meal intake amount estimation through deep learning. *Appl. Sci.*, 13(9):5742, 2023. 1
- [11] Filippos S. Konstantakopoulos, Eleni I. Georga, and Dimitrios I. Fotiadis. 3d reconstruction and volume estimation of food using stereo vision techniques. In *IEEE Int. Conf. Bioinform. Bioeng.*, pages 1–4. IEEE, 2021. 1
- [12] Minjae Kwak, Liang Mao, and Yujia Zheng. A cross-modal mutual knowledge distillation framework for alzheimer’s disease diagnosis: Addressing incomplete modalities. *IEEE Trans. Autom. Sci. Eng.*, 2025. 2
- [13] Fu Peng Wei Lo, Yifan Sun, Jianing Qiu, and Benny Lo. Food volume estimation based on deep learning view synthesis from a single depth map. *Nutrients*, 11(10):2447, 2019. 1
- [14] Jinge Ma, Xiaoyan Zhang, Gautham Vinod, Siddeshwar Raghavan, Jiangpeng He, and Fengqing Zhu. Mfp3d: Monocular food portion estimation leveraging 3d point clouds. *arXiv preprint arXiv:2411.10492*, 2024. 1, 2, 3, 4
- [15] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 2
- [16] Weiqing Shao, Wujie Min, Shengnan Hou, Mengjiang Luo, Ting Li, Yuanjie Zheng, and Shuqiang Jiang. Dpf-nutrition: Food nutrition estimation via depth prediction and fusion. *Foods*, 12(23):4293, 2023. 2
- [17] Zeman Shao, Shaobo Fang, Runyu Mao, Jiangpeng He, Jeannette L. Wright, Deborah A. Kerr, Carol J. Boushey, and Fengqing Zhu. Towards learning food portion from monocular images with cross-domain feature adaptation. In *IEEE Int. Workshop Multimedia Signal Process.*, pages 1–6. IEEE, 2021. 4
- [18] Zeman Shao, Gautham Vinod, Jiangpeng He, and Fengqing Zhu. An end-to-end food portion estimation framework based on shape reconstruction from monocular image. In *IEEE Int. Conf. Multimedia and Expo*, pages 942–947. IEEE, 2023. 4
- [19] Kai Sheng, Weijie Dong, Changsheng Ma, and Junliang Chen. A lightweight transformer-based deep neural network for food image recognition. *IEEE Access*, 10:42738–42749, 2022. 2
- [20] Gautham Vinod, Jiangpeng He, Zeman Shao, and Fengqing Zhu. Food portion estimation via 3d object scaling. In *CVPRW*, pages 3861–3870, 2024. 1, 2, 3, 4
- [21] Hao Wang, Yue Li, and Xinyu Zhang. Multiscale knowledge distillation with attention based fusion for robust human activity recognition. *Scientific Reports*, 2024. 2
- [22] Zhihui Wang, Xiaoming Chen, Yang Liu, and Kun Zhang. Visual nutrition analysis: Leveraging segmentation and regression for food nutrient estimation. *Front. Nutr.*, 11: 1505178, 2024. 2
- [23] Yuan Yue, Wenyan Jia, and Mingui Sun. Measurement of food volume based on single 2-d image without conventional camera calibration. In *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, pages 2166–2169. IEEE, 2012. 1
- [24] Yiming Zhang, Yang Liu, and Wei Chen. Cross-modal distillation for widely differing modalities. In *ICASSP*, 2024. 2