

# Transferring Knowledge for Food Image Segmentation using Transformers and Convolutions

Grant Sinha  
University of Waterloo  
Waterloo, Ontario, Canada  
gsinha@uwaterloo.ca

Amy Tai  
University of Waterloo  
Waterloo, Ontario, Canada  
amy.tai@uwaterloo.ca

Krish Parmar  
University of Waterloo  
Waterloo, Ontario, Canada  
k6parmar@uwaterloo.ca

Yuhao Chen  
University of Waterloo  
Waterloo, Ontario, Canada  
yuhao.chen1@uwaterloo.ca

Hilda Azimi  
National Research Council Canada  
Ottawa, Ontario, Canada  
hilda.azimi@nrc-cnrc.gc.ca

Alexander Wong  
University of Waterloo  
Waterloo, Ontario, Canada  
alexander.wong@uwaterloo.ca

Pengcheng Xi  
National Research Council Canada  
University of Waterloo, Ontario, Canada  
pengcheng.xi@nrc-cnrc.gc.ca

## Abstract

Food image segmentation is an important task that has ubiquitous applications, such as estimating the nutritional value of a plate of food. Although machine learning models have been used for segmentation in this domain, food images pose several challenges. One challenge is that food items can overlap and mix, making them difficult to distinguish. Another challenge is the degree of inter-class similarity and intra-class variability, which is caused by the varying preparation methods and dishes a food item may be served in. Additionally, class imbalance is an inevitable issue in food datasets. To address these issues, two models are trained and compared, one based on convolutional neural networks and the other on Bidirectional Encoder representation for Image Transformers (BEiT). The models are trained and evaluated using the FoodSeg103 dataset, which is identified as a robust benchmark for food image segmentation. The BEiT model outperforms the previous state-of-the-art model by achieving a mean intersection over union of 49.4 on FoodSeg103. This study provides insights into transferring knowledge using convolution and Transformer-based approaches in the food image domain.

## 1. Introduction

Every year, malnutrition causes a \$10 billion annual burden on the healthcare system in the United States [8]. There

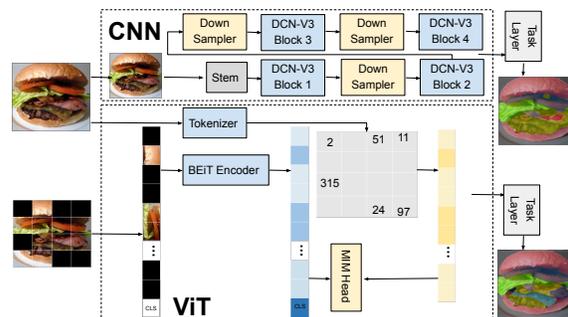


Figure 1. Comparison of downstream task performance of CNNs and ViTs on food image segmentation for the purpose of evaluating feature transferability.

is a strong link between poor nutrition and disease, mortality, and impaired quality of life [17, 24]. Older adults are at the greatest risk of health complications due to nutritional deficiencies, being over four times more likely to be hospitalized due to malnutrition. Further, they are at higher risk because one in four adults aged 65 years or older is malnourished [9, 10].

Nutrition monitoring has been proposed as a solution for combating malnutrition among the at-risk aging popu-

lation [15]. Methods exist for monitoring and measuring nutritional intake, but most of them are not reliable or accurate [18]. For example, food diaries, questionnaires, weighing, and photography each suffer from slowness, the need for trained personnel, or reporting biases.

Recognizing and breaking down the nutritional contents of a plate of food is one of the core tasks in food imaging for nutrition monitoring. Various pipelines using computer vision models have been designed [23, 28]. One promising approach [16] combines semantic segmentation with depth imaging to generate segmentation masks. The depth map is then used to compute estimates for food volumes.

Convolutional Neural Networks (CNN) are one of the mainstream methods for semantic segmentation. CNNs are mostly lightweight in terms of computational and memory requirements and benefit from strong inductive biases. Numerous variations of CNNs have been proposed, which introduce advanced architectural designs and complex convolutional operations. For example, deformable convolutions endow CNNs with adaptive receptive fields [5], and pyramid pooling modules allow CNNs to better make use of global context clues in the scene [29].

More recently, Transformer-based models have achieved great success across computer vision tasks including semantic segmentation [6]. Equipped with multi-headed self-attention mechanisms, these models have a global receptive field that enables them to make better use of global context information, although this comes at the expense of a higher computational and memory cost [25].

In addition to the ubiquity of Transformer-based models, recent years have also seen a shift towards large-scale vision models. BEiT [1, 14] and InternImage [27] are both representatives of this paradigm. They are large-scale models pre-trained on ImageNet and capable of being fine-tuned for numerous downstream vision tasks [7]. After pre-training, a task-specific layer or network is appended, and the whole network is fine-tuned for the desired downstream task.

Semantic segmentation on food images presents a unique set of challenges. One of them is that foods can occlude and mix with one-another in complex ways, preventing the model from accurately delimiting between them [15]. Another challenge is the presence of inter-class similarity and intra-class variability. For example, chicken and ham can be prepared to have a very similar appearance. Different cooking techniques applied to chicken can result in vastly different visual characteristics. An effective food segmentation model must be able to recognize the diverse appearances of chicken without confusing it with other similar-looking foods like ham [28].

Food image datasets are not as large or robust as other domains, especially when it comes to the need of image masks for training semantic segmentation models. UBIMIB2016 [4], UECFoodPixComplete [13], and Food-

Seg103 [28] are datasets with fine-grained segmentation masks. Other existing datasets are mainly for food recognition and coarse-grained segmentation tasks. For instance, Recipe1M comprises one-million food images with corresponding recipes [12], and Food101 contains approximately one-hundred thousand food images across 101 classes [11]. However, neither of them includes annotation masks for segmentations. Among the food datasets that support semantic segmentation, we identify FoodSeg103 as the most robust [28]. This is due to its detailed, pixel-level annotations, plus empirical evidence of its relatively high difficulty. Previous works have achieved a mean intersection over union (mIoU) over all classes of up to 45.1 using a ViT-B/16-based Segmentation Transformer [28, 30].

In this study, we aim to gain insight into how convolutional and Transformer-based architectures differ in knowledge transfer to the food image domain. Convolution-based methods are one of the main-stream approaches for segmentation, including the recent InternImage model [27], which is based on a deformable convolution operation, DCN-V3. It has achieved state-of-the-art results on the challenging ADE20k dataset for semantic segmentation [31]. Therefore, we include InternImage [27] as a baseline model.

In this work, we propose utilizing the BEiT v2 model [14], which is a recent Transformer-based model that has yet to be applied to food image segmentation. The BEiT v2 encoder is more sophisticated than prior works, owing to its pre-training procedure, which involves training an image tokenizer and visual codebook (see Figure 1). This promotes the learning of richer, semantic-level information, rather than lower-level, pixel-based information.

Through exploring the distinctions between InternImage and BEiT v2 on knowledge transfer to the domain of food images, we make the following findings and contributions:

- We survey the landscape of food image datasets for the purpose of semantic segmentation. We identify the most robust dataset and explain what makes it strong for semantic segmentation. Additionally, we explain why other datasets may not be suitable for this task.
- We train a tokenizer to investigate the power of vector-quantized knowledge distillation in the food image domain. We find that the tokenizer is learning semantic concepts in the food image dataset.
- We evaluate state-of-the-art models in the new domain of food images. Our experiments will provide insight into how knowledge transferability differs between convolution and Transformer-based architectures.
- We establish top results on the FoodSeg103 dataset. In particular, a BEiT v2 model achieves a mIoU of 49.4 on FoodSeg103, outperforming the previous state-of-the-art model’s mIoU of 45.1.

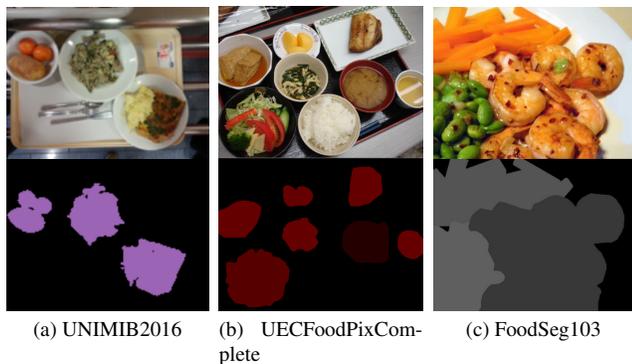


Figure 2. Image, mask pair from UNIMIB2016, UECFoodPixComplete, and FoodSeg103.

## 2. Methods

### 2.1. Dataset

The field of food image datasets has expanded in recent years, yet few datasets are specifically designed for semantic segmentation. Furthermore, each of these datasets has its own unique strengths and weaknesses. The datasets we investigate for semantic segmentation include UNIMIB2016, UECFoodPixComplete, and FoodSeg103 [4, 13, 28]. We choose these datasets because they all provide segmentation masks. Larger food image datasets exist but they do not contain segmentation masks. We include a sample image and mask pair from each of UNIMIB2016, UECFoodPixComplete, and FoodSeg103 in Figure 2.

UNIMIB2016 was one of the first datasets in the food domain with ground truth food masks to support food segmentation. It was published in 2017 by researchers from the University of Milano-Bicocca for the purpose of investigating food segmentation, food consumption, and dietary monitoring. The dataset consists of 1027 food images spanning 73 food categories. Food image segmentation masks are created manually with polygonal bounding boxes. Further, masks are not per-item, so all food classes are clubbed into one category in the segmentation masks [4].

UECFoodPixComplete was released in 2020 by researchers from the University of Electro-Communications. It includes 102 dishes and comprises 10,000 food images. The segmentation masks were obtained semi-automatically using GrabCut, which segments images based on user-initialized seeding [21]. The automatically-generated masks were further refined by human annotators based on a set of predefined rules [13].

FoodSeg103 is a recent dataset designed for food image segmentation, consisting of 7,118 images depicting 730 dishes. It was released in 2021 and the dataset’s masks are pixel-level. They were obtained through manual annotations. The masks were then inspected for further refine-

Table 1. A summary of several robust food datasets, their number of classes, number of images, and whether it includes masks for segmentation ground truths.

Dataset	Classes	Images	Masks
UNIMIB2016	73	1027	Yes
UECFoodPixComplete	102	10,000	Yes
FoodSeg103	103	7118	Yes
Recipe1M		887,706	No
Food101	101	101,000	No

ment. Compared to UECFoodPixComplete, FoodSeg103 proves to be a more challenging benchmark for food image segmentation. In experiments conducted with a deeplabv3+ model [2], a lower mIoU of 34.2 was achieved on FoodSeg103 than UECFoodPixComplete [28]. Further, unlike UECFoodPixComplete, which covers entire dishes but lacks fine-grained annotation for individual dish components, FoodSeg103 aims to annotate dishes at a more fine-grained level, capturing the characteristics of each dish’s individual ingredients.

We conclude that FoodSeg103 is the most suitable dataset for training and evaluating a model for food image segmentation. In Table 1, we summarize various datasets, including some larger datasets that lack annotation masks, which are essential for training segmentation models using a supervised paradigm.

### 2.2. InternImage

Neural networks that make use of convolutions have been a longstanding standard in computer vision. The new InternImage model, which uses a variation of convolutions called Deformable Convolution V3 (DCN-V3), has achieved state-of-the-art results in semantic segmentation on notable benchmarks such as ADE20k [31]. The DCN-V3 operation uses a  $3 \times 3$  convolution kernel with learnable receptive fields and modulation scalars [27]. By equipping convolutions with a learnable receptive field, they can overcome their weaknesses. In particular, they can learn to utilise long-range dependencies and adaptive spatial aggregation from the data. Further, InternImage also retains the benefits of computational and memory efficiency that most convolutional models have over Transformer-based models.

InternImage processes images in several stages using blocks (see Figure 1). The input is down-scaled between the blocks so that each feature map has a different resolution. Each DCN-V3 block consists of a number of basic blocks, where the basic block is designed to resemble a Vision Transformer, at a high level. The basic block consists of two sub-layers. The first uses the DCN-V3 operation as its core operator, and the second is a simple feedforward network. Both outputs have layer normalization applied.

Table 2. A summary of the model versions considered in experiments, including their sizes and input resolutions.

Model Name	Number of Parameters	Crop Size
BEiT Large	441M	640×640
BEiT v2 Base	163M	512×512
BEiT v2 Large	441M	512×512
BEiT-3	1B	896×896
InternImage-B	128M	512×512
InternImage-L	256M	640×640
InternImage-XL	368M	640×640
InternImage-H	1.31B	896×896

Finally, each sub-layer has a residual connection surrounding it [27]. InternImage is trained on ADE20k and equipped with Mask2Former (a unified framework for all segmentation tasks based on masked attention) [3]. For the task layer, an UperNet decoder is appended for semantic segmentation.

### 2.3. BEiT

The BEiT models [1, 14, 26] are Vision Transformer-based architectures capable of encoding images. They use masked image modeling as the pre-training task. In this process, images with corrupted patches are fed into a vision encoder, with the reconstruction target being visual tokens obtained from an image tokenizer.

There are three versions of the BEiT model. The first version introduces masked image modeling as a pre-training task for ViTs, using a pre-trained tokenizer from DALL-E [14, 20]. The second version improves upon the first by training its own visual tokenizer using ImageNet-1K [1]. Finally, the third version is multi-modal and significantly scales up the number of parameters (1B parameters used on vision tasks) [26].

We chose to use BEiT v2 for several reasons. Although BEiT v3 achieves stronger results via scale and multi-modality, its approach to vision is not upgraded from BEiT v2. We also selected BEiT v2 over the first BEiT due to its learned tokenizer, which provides a richer reconstruction target for the encoder than the pre-trained tokenizer used in the first BEiT model. The BEiT v2 model uses a teacher model to train a tokenizer using vector-quantized knowledge distillation [1]. The tokenizer consists of a Transformer encoder and quantizer. The quantizer works by mapping each vector output from the Transformer encoder to its nearest neighbour in a codebook, which serves as a visual vocabulary. In this way, the continuous semantic space that the Transformer maps to is quantized into discrete codes. A decoder then learns to construct teacher model outputs from the sequence of codebook vectors. Through this, the tokenizer learns a meaningful codebook of semantic con-

Table 3. An overview of the models applied to the FoodSeg103 dataset, along with their mean intersection over union. InternImage and CCNet are based on convolutions; the remainder are on Transformers

Model Name	Number of Parameters	mIoU
<b>BEiT v2 Large</b>	<b>441M</b>	<b>49.4</b>
InternImage-B	128M	41.1
SeTR-MLA	711M	45.1
SeTR-Naive	723M	43.9
Swin-S	931M	41.6
CCNet	381M	35.5

cepts from the training set. This is an improvement over the first version of BEiT, which used a pretrained tokenizer from DALL-E [20].

The aim of the tokenizer is to discretize the continuous semantic space into codebook categories. Discretizing the continuous semantic space into compact codes using the codebook allows us to create a semantic-aware tokenizer. Each code in the codebook represents a semantic concept from the dataset used (e.g., unique codes for plates, eyes, foods, etc.). This approach will promote the mask-image modeling from pixel to semantic-level learning, resulting in richer representations. In the food domain, where the same ingredient can have varied appearances and contexts in different images, these stronger representations will prove to be useful.

After training the tokenizer, a network of stacked Transformer blocks is trained as a vision encoder. The encoder is trained using the task of masked image modeling. During this process, the input image is split into patches, and some of the patches are masked using a learnable embedding. This corrupted image is then fed as a sequence into the Transformer blocks, with the goal of recovering visual codebook tokens from the input image. The tokenizer trained in the previous step is used to generate the ground truth values used by the BEiT v2 encoder in this step. Once the encoder has been trained, we append an UperNet decoder as the task layer for semantic segmentation, then fine-tune the model.

## 3. Results

### 3.1. InternImage Model

Among several sizes of the InternImage model that we considered for experiments, we use the base InternImage-B model to compare against the large BEiT v2 model. They use the same crop size and have a comparable number of parameters. The relative sizes of the various InternImage and BEiT models are summarized in Table 2.

The InternImage encoder is pre-trained before having



Figure 3. Food images to be tokenized using the learned tokenizer and codebook. Same-category foods have the highest IoUs, but similarly-looking cross-category foods can also yield high IoUs.

Table 4. Token intersection over unions for images in Figure 3

	Prime Rib I	Prime Rib II	Filet Mignon I	Filet Mignon II	Pork Chop I
Prime Rib I	-	<b>0.281</b>	0.121	0.153	0.131
Prime Rib II	0.281	-	0.123	0.126	0.152
Filet Mignon I	0.121	0.123	-	<b>0.242</b>	<b>0.233</b>
Filet Mignon II	0.153	0.126	0.242	-	0.171
Pork Chop I	0.131	0.152	0.233	0.171	-

a task-specific layer appended for semantic segmentation. pre-training for InternImage-B is done using classification on ImageNet-1K for 300 epochs. [27].

Following the pre-training, the InternImage models are fine-tuned for semantic segmentation on FoodSeg103. In particular, we append an UperNet decoder as a task layer for semantic segmentation. Subsequently, the model undergoes end-to-end fine-tuning using the AdamW optimizer for 160K iterations. We use a learning rate of  $6e-05$ , no weight decay, and betas of (0.9, 0.999).

## 3.2. BEiT Model

### 3.2.1 BEiT Tokenizer Training

We trained a BEiT v2 tokenizer on Food101 to investigate the power of vector-quantized knowledge distillation. The visual tokenizer is initialized as a ViT-B/16 Transformer and taught by a CLIP-B/16 model [19]. The codebook size is 8192, and the dataset used is Food101, which has 101,000 images. The model is trained for 100 epochs, and the input image resolution is  $224 \times 224$ , which gives  $14 \times 14$  total image patches each with resolution  $16 \times 16$ .

To evaluate the trained tokenizer, we run inference on food images from several categories in Food101, obtaining their corresponding token sequence. We compare tokenized images by considering the IoU of their sets of tokens. In particular, given tokenized images  $\mathcal{I}_i, \mathcal{I}_j \subseteq \{1, \dots, 8196\}$ , we compute

$$\frac{|\mathcal{I}_i \cap \mathcal{I}_j|}{|\mathcal{I}_i \cup \mathcal{I}_j|}$$

to measure the similarity of the tokenized images. We find that images belonging to the same food category typically have a higher degree of similarity, indicating that the tokenizer is successfully learning semantic concepts for the codebook. Figure 3 and Table 4 show a collection of food images and their corresponding token IoUs. Encouragingly, we see that the highest IoUs come from foods of the same class. However, we also see the inter-class similarity of the food domain in the similarity of the tokenized pork chop and filet mignon.

### 3.2.2 BEiT model tuning

After learning the visual tokenizer, pre-training of the BEiT v2 encoder is conducted. For the large-size model, it is initialized as a ViT-L/16 large-size Vision Transformer with a 16-pixel patch size. The tokenizer that generates ground truths for the pre-training of masked image modeling is used here. Masked image modeling is performed using ImageNet-1k at a resolution of  $224 \times 224$  for 1600 epochs, resulting in  $14 \times 14$  image patches. A masking ratio of 40% is used, which means that up to 75 patches will be masked per image.

Given the pre-trained large-size BEiT v2 encoder, we append an UperNet decoder for the task layer. The entire model is then fine-tuned for 160K iterations (approximately 50 epochs) at an input resolution  $512 \times 512$  using the AdamW optimizer. We use a learning rate of  $3e-05$ , weight decay of 0.05, and betas of (0.9, 0.999). The best-performing model during training achieved a mIoU of 49.4

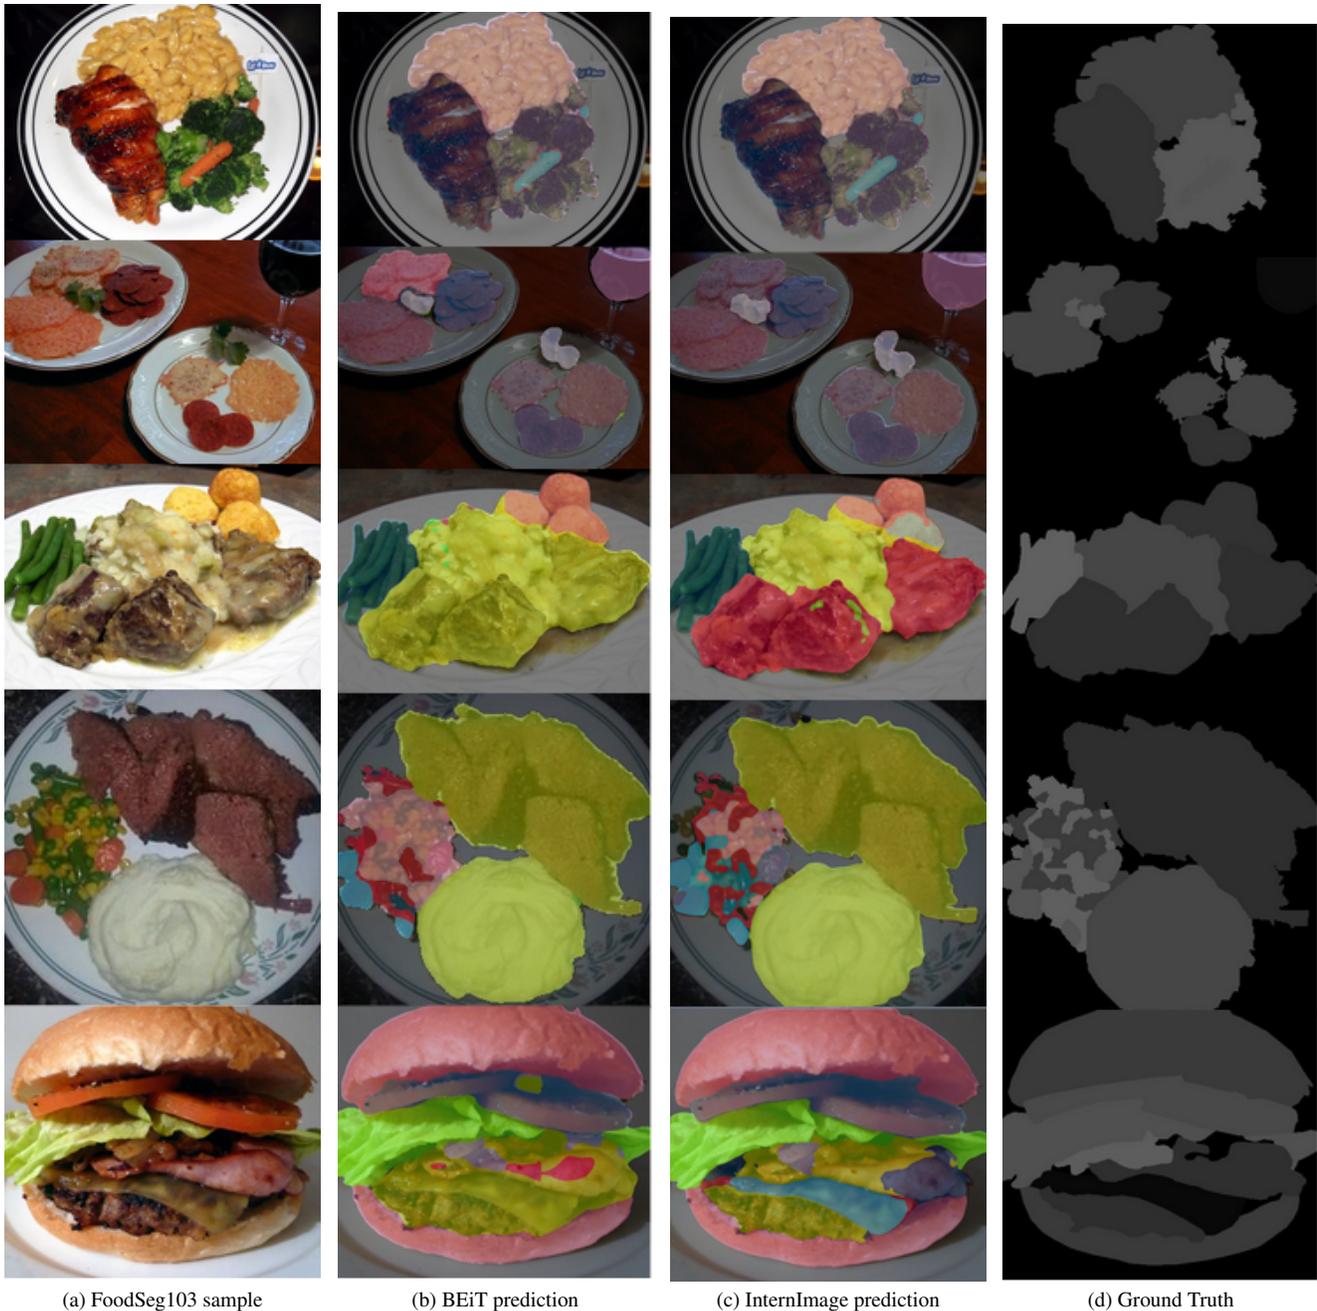


Figure 4. A comparison of inference on FoodSeg103 by the best-performing BEiT v2 and InternImage-B model.

after 160k iterations.

### 3.3. Performance of Models on FoodSeg103

For the BEiT v2 model, during training, the best-performing model achieved a mIoU of 49.4 at iteration 160K. This is the strongest result reported on FoodSeg103, representing a new state-of-the-art for food segmentation on the dataset. On the other hand, for the InternImage-B

model, during training, the best-performing model achieved a mIoU of 41.1 at iteration 48k. Table 3 shows the performance of these models in contrast to the other models previously applied to FoodSeg103. BEiT v2 Large outperforms the prior state-of-the-art model and uses less parameters.

Figure 4 shows a visual comparison of segmentation mask predictions made by our trained models. Although both models make good segmentation prediction masks,

these randomly selected testing cases prove to be challenging in terms of detecting food items. For example, on second row, the InternImage model failed to segment the food item on the top left corner. In the third example, the BEiT model failed to segment meat chunks from neighboring food items. In the last example, both models were only able to segment parts of the bacon in the burger.

## 4. Discussions

Several challenges affect the performance of the model in food image segmentation. The distribution of ingredients in food image segmentation is long-tailed, resulting in sparse data for ingredients in the long-tail, and thus, the model poorly predicts these categories. In the FoodSeg103 dataset used, we identified several problematic foods in the long-tail that the model struggled to segment correctly. hamburger appeared 7 times in the training set and 1 time in the test set, pudding appeared 5 times in the train set and 1 time in the test set, and kelp appeared 4 times in the train set and 5 times in the test set. These foods posed difficulty for each model. Figure 4 shows an example of a hamburger, which both models struggle to accurately segment.

A second challenge unique to food is the degree of inter-class similarity and intra-class variability. Each class in food images represents a food ingredient that can be prepared in numerous ways, resulting in a wide variation in appearance. This can also lead to visually similar-looking foods being mistaken for each other. This can be seen in Figure 3 and Table 4, where a visually similar filet mignon and pork chop are tokenized similarly. The models must be capable of understanding that a food ingredient prepared in any way is still the same food, while also differentiating between similar-looking foods [15, 16, 28].

Owing to the superior abilities of Transformers, BEiT exhibited strong improvement over the state-of-the-art, while the performance of InternImage plateaued at a lower mIoU. One factor we can attribute this to is the global receptive field of Transformers. Even though InternImage uses learnable deformable convolutions, its kernel size is still restricted to  $3 \times 3$ , which prevents InternImage from having the same level of global understanding as BEiT. Global context undoubtedly gives valuable insight into categorizing pixels, as certain food items are more common to occur alongside other specific food items. Another factor we may attribute BEiT v2's performance to is the use of vector-quantized knowledge distillation to train a tokenizer for BEiT's pretraining reconstruction targets. This pretraining routine endows BEiT v2 with a stronger grasp of the dataset semantics than InternImage, as the continuous, high-dimensional semantic space is quantized into codes, enabling the model to learn a visual vocabulary through its codebook.

Although BEiT outperformed InternImage in mIoU,

there are instances where BEiT struggles to distinguish between classes where InternImage does not. For instance, in the third sample of Figure 4(b), BEiT groups the mashed potato and meat together, whereas InternImage was able to differentiate between them.

## 5. Conclusion

In this work, we sought to evaluate the knowledge transfer capability of Transformer and convolution-based vision backbones on the downstream task of food image segmentation. We select BEiT v2 and InternImage as strong representatives for Transformer and convolutional approaches respectively. Using these representatives and the benchmark FoodSeg103, we found that Vision Transformers have superior downstream task transferability to convolutional networks for the task of food image segmentation.

For future direction, pre-training BEiT v2 on Food101 could obtain stronger representations for food images than if it were trained on ImageNet-1K [1, 22]. Furthermore, the multi-modal capabilities of BEiT-3 could be used to obtain even stronger food representations [26].

## References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image transformers. 2, 4, 7
- [2] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 3
- [3] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 4
- [4] Gianluigi Ciocca, Paolo Napoletano, and Raimondo Schettini. Food recognition: A new dataset, experiments, and results. 21(3):588–598. 2, 3
- [5] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. 2
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. 2
- [7] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggong Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. EVA: Exploring the limits of masked visual representation learning at scale. 2
- [8] Scott Goates, Kristy Du, Carol A. Braunschweig, and Mary Beth Arensberg. Economic burden of disease-associated malnutrition at the state level. 11(9):e0161833. 1
- [9] Matthias J. Kaiser, Jürgen M. Bauer, Christiane Rämisch, Wolfgang Uter, Yves Guigoz, Tommy Cederholm, David R.

- Thomas, Patricia S. Anthony, Karen E. Charlton, Marcello Maggio, Alan C. Tsai, Bruno Vellas, Cornel C. Sieber, and Mini Nutritional Assessment International Group. Frequency of malnutrition in older adults: a multinational perspective using the mini nutritional assessment. 58(9):1734–1738. 1
- [10] David P. Lanctin, Francheska Merced-Nieves, Renee M. Mallett, Mary Beth Arensberg, Peggi Guenter, Suela Sulo, and Timothy F. Platts-Mills. Prevalence and economic burden of malnutrition diagnosis among patients presenting to united states emergency departments. 28(3):325–335. 1
- [11] Bossard Lukas, Guillaumin Matthieu, and Luc Van Gool. Food-101 - mining discriminative components with random forests. 2
- [12] Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusaf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. 2
- [13] Kaimu Okamoto and Keiji Yanai. UEC-FoodPix complete: A large-scale food image segmentation dataset. In Alberto Del Bimbo, Rita Cucchiara, Stan Sclaroff, Giovanni Maria Farinella, Tao Mei, Marco Bertini, Hugo Jair Escalante, and Roberto Vezzani, editors, *Pattern Recognition. ICPR International Workshops and Challenges*, volume 12665, pages 647–659. Springer International Publishing. Series Title: Lecture Notes in Computer Science. 2, 3
- [14] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. BEIT v2: Masked image modeling with vector-quantized visual tokenizers. 2, 4
- [15] Kaylen J. Pfisterer, Robert Amelard, Jennifer Boger, Audrey G. Chung, Heather H. Keller, and Alexander Wong. Enhancing food intake tracking in long-term care with automated food imaging and nutrient intake tracking (AFINI-t) technology. 2, 7
- [16] Kaylen J. Pfisterer, Robert Amelard, Audrey G. Chung, Braeden Syrnyk, Alexander MacLean, Heather H. Keller, and Alexander Wong. Automated food intake tracking requires depth-refined semantic segmentation to rectify visual-volume discordance in long-term care homes. 12(1). 2, 7
- [17] M. Pirlich and H. Lochs. Nutrition in the elderly. 15(6):869–884. 1
- [18] Ross L. Prentice, Yasmin Mossavar-Rahmani, Ying Huang, Linda Van Horn, Shirley A. A. Beresford, Bette Caan, Lesley Tinker, Dale Schoeller, Sheila Bingham, Charles B. Eaton, Cynthia Thomson, Karen C. Johnson, Judy Ockene, Gloria Sarto, Gerardo Heiss, and Marian L. Neuhouser. Evaluation and comparison of food records, recalls, and frequencies for energy and protein assessment by using recovery biomarkers. 174(5):591–603. 2
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. 5
- [20] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. 4
- [21] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "GrabCut": interactive foreground extraction using iterated graph cuts. In *ACM SIGGRAPH 2004 Papers*, pages 309–314. ACM. 3
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. 7
- [23] Quin Thames, Arjun Karpur, Wade Norris, Fangting Xia, Liviu Panait, Tobias Weyand, and Jack Sim. Nutrition5k: Towards automatic nutritional understanding of generic food. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. ISSN: 10636919. 2
- [24] D. R. Thomas, W. Ashmen, J. E. Morley, and W. J. Evans. Nutritional management in long-term care: development of a clinical guideline. council for nutritional strategies in long-term care. 55(12):M725–734. 1
- [25] Ashish Vaswani, Google Brain, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2
- [26] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a foreign language: BEIT pretraining for all vision and vision-language tasks. 4, 7
- [27] Wenhui Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, Xiaogang Wang, and Yu Qiao. InternImage: Exploring large-scale vision foundation models with deformable convolutions. 2, 3, 4, 5
- [28] Xiongwei Wu, Xin Fu, Ying Liu, Ee-Peng Lim, Steven C. H. Hoi, and Qianru Sun. A large-scale benchmark for food image segmentation. 2, 3, 7
- [29] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. 2
- [30] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H S Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. 2
- [31] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20k dataset. 2, 3