

In The Wild Ellipse Parameter Estimation for Circular Dining Plates and Bowls

Akil Pathirana^{1*}, Chris Czarnecki^{1*}, Yuhao Chen^{1†}, Pengcheng Xi², Linlin Xu¹, Alexander Wong¹

¹ Vision and Image Processing Lab, University of Waterloo, ² National Research Council Canada

{alkarago, cczarnec, yuhao.chen1}@uwaterloo.ca

pengcheng.xi@nrc-cnrc.gc.ca

{l44xu, alexander.wong}@uwaterloo.ca

Abstract

Ellipse estimation is an important topic in food image processing because it can be leveraged to parameterize plates and bowls, which in turn can be used to estimate camera view angles and food portion sizes. Automatically detecting the elliptical rim of plates and bowls and estimating their ellipse parameters for data in-the-wild is challenging: diverse camera angles and plate shapes could have been used for capture, noisy background, multiple non-uniform plates and bowls in the image could be present. Recent advancements in foundational models offer promising capabilities for zero-shot semantic understanding and object segmentation. However, the output mask boundaries for plates and bowls generated by these models often lack consistency and precision compared to traditional ellipse fitting methods. In this paper, we combine ellipse fitting with semantic information extracted by zero-shot foundational models and propose WildEllipseFit, a method to detect and estimate the elliptical rim for plate and bowl. Evaluation on the proposed Yummy-ellipse dataset demonstrates its efficacy and zero-shot capability in real-world scenarios.



Figure 1. (a): the input image containing soup in a bowl. (b): the bowl segmented using GroundedSAM [14], where the elliptical rim has been lost. (c): candidate ellipses generated by our method. (d): final prediction (red) and the ground truth ellipse (green).

1. Introduction

In the realm of food image processing, accurately estimating ellipses has posed a persistent challenge, drawing attention for a decade [7, 16]. Food items are commonly presented in plates, bowls, or cups, and when viewed from various angles, the rims of these containers appear as ellipses. Leveraging this elliptical representation, we can estimate both the camera perspective and the portion size relative to the original container [5–7]. This characteristic renders ellipse estimation an ideal technique for two important downstream applications. Firstly, estimated ellipses serve as reference points for quantifying the contents of food items, aiding in the assessment of dietary intake. Secondly, they facilitate the estimation of rich multimodal information essential for finely controlled generation of food images or videos, such as camera angles or reference portion sizes. The success of both applications requires good method generalization to in-the-wild data which is characterized by complex food types, shapes and arrangements alongside plates, bowls, and utensils.

Several ellipse fitting methods exist such as the Direct Least Squares method [4], posing the ellipse fitting problem as an unconstrained optimization problem. Current general-purpose state-of-the-art method for ellipse fitting in-the-wild is that of Meng *et al.*, AAMED [11], leveraging edge detection followed by elliptic arc segmentation and a subsequent search of an arc adjacency matrix. Both traditional and data-driven methods for ellipse fitting lack a profound understanding of the semantic context of objects within the image [9, 11]. Although segmentation-based foundational models [8] have excelled in general object segmentation tasks, they encounter difficulties in accurately delineating plate and bowl segmentation masks with precise boundaries for ellipse estimation.

In this paper, we propose WildEllipseFit, a framework to automatically detect and estimate the elliptical parameters describing the rims of circular bowls and dining plates in an image. The framework combines Canny edge detection

*authors contributed equally

†corresponding author

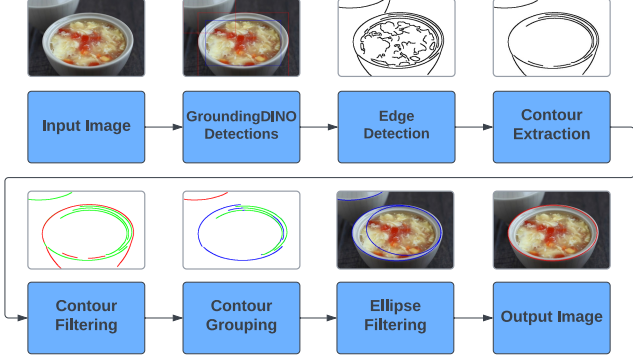


Figure 2. A block diagram for the ellipse fitting process with an example image.

[2] and contour processing along with object detection from GroundingDINO [10] to create accurate ellipse predictions. We also introduce an innovative contour filtering procedure, yielding good results on in-the-wild plate ellipse fitting.

To facilitate evaluation of our method we also introduce the Yummly-ellipse dataset. We have annotated images from the Yummly dataset [12] with manually-drawn ellipses around the rims of plates and bowls in the scene. The Yummly-ellipse dataset constitutes our second contribution towards more accurate plate ellipse fitting making it a good candidate for a benchmark dataset for the plate-ellipse-fitting problem.

2. Method

2.1. Method Overview

The overall process is visualized in Figure 2. First, GroundingDINO [10] is used to obtain bounding box detections for food and plates in the image. Then Canny edge detection [2] is applied to the image resulting in an edge map. Subsequently, curved contours that could belong to an ellipse are extracted and then filtered to remove contours not belonging to the rim. After all irrelevant contours have been filtered, the contours are grouped to determine which belong to the same ellipse. Once all contours have been grouped, ellipse fitting is applied to each group of contours to generate ellipse predictions. These prediction ellipses are finally filtered using semantic information from GroundingDINO.

2.2. GroundingDINO Detections & Edge Detection

The process used to localize ellipses leverages both GroundingDINO [10] and contour processing to determine whether ellipse predictions are correct. GroundingDINO is an open-set object detector capable of detecting arbitrary objects. For each image I , it is used to obtain bounding boxes around plates, bowls, and food. Food bounding boxes are combined into a larger bounding box under two conditions: 1) when they are contained within the bounding box of the same

plate/bowl or, 2) when they meet an Intersection-over-Union (IoU) threshold, G_M . This is done to cover the entirety of the food items on a given dish with one collective bounding box. Subsequently, the image is processed using the Canny edge detection [2] algorithm to obtain an edge map \mathcal{P} . What follows is extraction of elliptical edges.

2.3. Contour Extraction

Contours from the edge map \mathcal{P} are obtained, with each contour P being a sequence of pixels:

$$P = ((x, y))_{x \in [0;w], y \in [0;h]} \quad (1)$$

Where h is the height of the image I and w is its width. To separate elliptical edges of plates and bowls from jagged edges of food, information about their curvature can be used. We employ a similar method to that of Prasad [13].

Elliptical edges within the image will not have large changes in the tangential slope when compared to the jagged edges in the image. If each contour P was represented as a non-discrete curve, $f(x, y)$, the second derivative of $f(x, y)$ $\frac{d^2 f}{dx^2 dy^2}$ should be low for elliptical contours. Since points in the sequence P represent discrete pixels, the smallest possible difference in either x or y is 1, which means the concept of a derivative cannot directly be applied. We use vectors to represent the change in x and the change in y between two points in the sequence. Let \mathbf{d}_{p_u, p_v} be the vector representing the change in x and y between two pixels, where $p_u, p_v \in P$ and p_u, p_v are separated in the sequence by s steps, i.e. $v = u + s$. s is a hyperparameter.

Hence:

$$\mathbf{d}_{p_u, p_v} = \mathbf{p}_v - \mathbf{p}_u \quad (2a)$$

$$= \begin{bmatrix} x_v \\ y_v \end{bmatrix} - \begin{bmatrix} x_u \\ y_u \end{bmatrix} \quad (2b)$$

Subsequently, we consider the L1 distance, a , between consecutive \mathbf{d} vectors: If a is less than or equal to a set threshold value, ϵ , then a candidate contour $C = \emptyset$ is updated to the union of C and the subsequence $[p_v; p_{v+s}]$ from P :

$$C \leftarrow \{C \cup [p_v; p_{v+s}] : a \leq \epsilon\} \quad (3)$$

If the a is greater than ϵ , then the points in the aforementioned interval of the sequence P are discarded. The current set C of accepted points is checked to see if its length is at least L . If so it is recorded as a new contour:

$$\mathcal{C} \leftarrow \{\mathcal{C} \cup C : |C| \geq L\} \quad (4)$$

The set of accepted points C is then reassigned as $C = \emptyset$. This is to avoid any gaps or duplicate points in the subsequently captured contours. This process is repeated for every edge $P \in \mathcal{P}$. Ultimately, a set of contours \mathcal{C} is obtained representing only the edges with low curvature.




Figure 3. Example of curve extraction process, top left is original image. Middle is the Canny edge detection algorithm applied to the entire image. The right image is the contours extracted as a result of applying the curve extraction process.

2.4. Contour Filtering

As a result of the contour extraction algorithm, straight contours are also preserved. Thus, the next step is to remove them. Two extremal points of each contour $C \in \mathfrak{C}$ are identified and a line section (chord) is drawn between the two points. The chord can be represented as a vector χ .

The smallest possible distance between each point in the contour $p \in C$ and χ is then recorded. This is equivalent to taking the Euclidean distance between \mathbf{p} and χ . Then we remove points from \mathfrak{C} based on the following condition:

$$\mathfrak{C} \leftarrow \left\{ \mathfrak{C} \setminus C : \max_{\mathbf{p} \in C} \|\mathbf{p} - \chi\|_2^2 < d \right\} \quad (5)$$

This process is repeated $\forall C \in \mathfrak{C}$.

The final step in Contour Filtering is used to focus only on the top-surface elliptical contours of plates and bowls. When viewed from the side, bowls often have two elliptical contours, the one representing the top surface delineated by the bowl rim and one that can be observed as a result of viewing it from the side, stemming from the natural curvature of a bowl dish. Thus, it becomes necessary to distinguish which contour describes the rim of the bowl.

For every food bounding box predicted by GroundingDINO and its closest plate/bowl bounding box, there should be a gap between the bottom edges of the bounding boxes. Only contours and their points within this gap are considered. The y-distance between the bottom edge of the food bounding box and the lowest point in the Y coordinate for each contour is calculated and added to the set \mathcal{D} :

$$\mathcal{D} \leftarrow \{ |y_{\text{bbox},f} - y_{\text{min},C}| : \forall C \in \mathfrak{C} \} \quad (6)$$

Then the set of contours \mathfrak{C} and set of y-distances \mathcal{D} are updated iteratively:

$$\begin{aligned} &\textbf{while } |\max \mathcal{D} - \min \mathcal{D}| > H \textbf{ do} \\ &\quad U \leftarrow \left\{ \left| d - \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} d_i \right| : \forall d \in \mathcal{D} \right\} \\ &\quad \mathfrak{C} \leftarrow \{ \mathfrak{C} \setminus C_j : j = \arg \max U \} \\ &\quad \mathcal{D} \leftarrow \{ \mathcal{D} \setminus d_k : k = \arg \max U \} \\ &\textbf{done} \end{aligned} \quad (7)$$

Iteratively discarding contours of which y-distance to the bottom of the food instance bounding box deviates the most from the mean of all such distances allows us to filter out outlier contours. If $|\mathcal{D}| = 2$ and $|\max \mathcal{D} - \min \mathcal{D}|$ is still $> H$, we reject the contour for which $d \in \mathcal{D}$ is the largest.

This method is more robust than just removing the lowest-positioned contours because for some samples the curve extraction process does not remove some of the food instance edges. Rejecting the lowest-positioned contours would inevitably eliminate all valid bowl contours under such circumstances.

2.5. Contour Grouping

Now, having obtained the individual contours needed, the contours of the same ellipse must be combined. To check if two contours, C_1 and C_2 , are to be combined, we combine their sets of points into one set, $C_{C_1 \cup C_2} = C_1 \cup C_2$. If fitting an ellipse to C results in a close fit to the points, then they are likely to be a part of the same ellipse. For every point in the combined set, we get the squared Euclidean distance to the closest point on the fitted ellipse E . The mean of these distances is then taken as a score S for ellipse fitting:

$$S = \frac{1}{|C_{C_1 \cup C_2}|} \sum_{\mathbf{p} \in C_{C_1 \cup C_2}} \|\mathbf{p} - E\|_2^2 \quad (8)$$

If the score is lower than a set limit, M , the two contours are combined and their union replaces the two contours in the set \mathfrak{C} :

$$\mathfrak{C} \leftarrow \{ (\mathfrak{C} \setminus \{C_1, C_2\}) \cup \{C_1 \cup C_2\} : S < M \} \quad (9)$$

2.6. Ellipse Filtering

After the contour merging algorithm, a convex hull of each contour is first found and then fitted with OpenCV's default ellipse fitting implementation [1, 3], which we further denote as operator name: "fit". We leverage the bounding box information from GroundingDINO to filter out incorrect ellipses. Each ellipse is first compared with the closest plate/bowl bounding box. Intuitively, a well-fitted ellipse prediction should not have a large portion of its area outside of the bounding box. This idea can be used to filter out ellipses that

	G_M	s	ϵ	L	d	H	M	A_p	D_f
A	0.6	8	2	60	7.0	10	150	0.08	450
B	0.7	7	2	60	7.5	10	125	0.06	450

Table 1. Optimal Parameter sets for both method A and B calculated on training dataset of 300 images.

have a minimum area fraction, A_p , outside the bowl or plate bounding box. A good ellipse prediction would also have its center coordinates, (x_e, y_e) , close to the center of any of the food bounding boxes. Therefore, the distance between the center of each ellipse and the closest food bounding box can also be thresholded by a maximum distance, D_f :

$$\mathcal{E} \leftarrow \left\{ \mathcal{E} \setminus \{E\} : \frac{\text{Area}(E \setminus \text{bbox_p})}{\text{Area}(E)} \geq A_p \right\} \quad (10a)$$

$$\mathcal{E} \leftarrow \left\{ \mathcal{E} \setminus \{E\} : \|\mathbf{e}_c - \mathbf{b}_c\|_2^2 \geq D_f \right\} \quad (10b)$$

Where bbox_p stands for the plate bounding box, $\mathbf{e}_c = [x_e \ y_e]^T$ and $\mathbf{b}_c = [x_{\text{bbox_f_center}} \ y_{\text{bbox_f_center}}]^T$, where bbox_f is the food bounding box.



Figure 4. Plate filtering.




Figure 5. Food distance filtering.

After these two filters have been applied, the remaining ellipse prediction parameters are returned.

3. Results

We applied our algorithm to a collection of 400 images from the Yummly-28K dataset [12], which contains a wide variety of food scenes under many different viewing angles. All 400 images contained at least one elliptical contour and were manually annotated with ground truth ellipses. We name the resulting dataset Yummly-ellipse.

We have chosen to use the Chamfer distance between the perimeters of the predicted ellipse and the nearest ground truth ellipse as the metric for evaluation as it incorporates all errors of the ellipse parameters under one metric. The Chamfer distance between two sets of points, A and B , is defined as:




Figure 6. Sample correct detections by our framework.



Figure 7. Sample incorrect detections by our framework.

$$D_{\text{Chamfer}}(A, B) = \frac{1}{2} \left(\sum_{x \in A} \min_{y \in B} d(x, y) + \sum_{y \in B} \min_{x \in A} d(x, y) \right) \quad (11)$$

In our case we adopt Euclidean distance for $d(x, y)$. We choose two methods of comparing ground truth ellipses to predictions. The first method, method A, takes each ground truth ellipse and calculates the Chamfer distance between its perimeter and the nearest predicted ellipse’s perimeter. The second method, method B, takes each predicted ellipse and calculates the Chamfer distance between its perimeter and the nearest ground truth ellipse’s perimeter. Method A penalizes missed ellipse predictions, assigning more importance to recall. Method B focuses more on the precision of each individual ellipse prediction, which is an important consideration if WildEllipseFit is to be used for camera angle estimation.

The threshold for detection used by GroundingDINO for the plates, bowls, and food was set at 0.35 with prompts of "Plate", "Bowl", and "Food". For Canny Edge detection, we used Scikit-Image’s [15] Canny edge detection implementation with a σ of 2.5.

Method	μ_A	σ_A	N	μ_B	σ_B	N
AAMED	19.33	32.57	68	38.91	56.05	68
Ours	16.87	33.36	37	7.11	7.61	35

Table 2. Baseline Chamfer Distance metric for the method proposed by Meng *et al.* [11] versus WildEllipseFit on a test dataset of 100 images. μ_A indicates Chamfer distance calculated according to Method A, μ_B according to Method B. σ indicates standard deviation. N refers to number of images with atleast one ellipse prediction. Optimal mean Chamfer Distance has been bolded.

We compare our method in Table 2 using the optimal parameters found previously with the arc adjacency matrix method proposed by Meng *et al.* [11] using optimal hyperparameters reported by the authors. We find that while our method results in fewer images with ellipse predictions, we

obtain a lower Chamfer distance indicating a higher precision.

4. Conclusion

In this paper, we discussed the motivation for estimating elliptical parameters of circular dining plates and bowls. We outlined WildEllipseFit, a multi-stage process combining contour processing with semantic information from zero-shot models to estimate ellipse parameters. We demonstrate its performance on a dataset of 400 images annotated with ground truth ellipses. The proposed framework was compared with current ellipse-fitting methods and the improvements provided by incorporating zero-shot models was shown.

References

- [1] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. 3
- [2] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, 1986. 2
- [3] Andrew Fitzgibbon and Robert Fisher. A buyer's guide to conic fitting. 1970. 3
- [4] A. Fitzgibbon, M. Pilu, and R.B. Fisher. Direct least square fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):476–480, 1999. 1
- [5] Wenyan Jia, Yaofeng Yue, John D. Fernstrom, Ning Yao, Robert J. Scabassi, Madelyn H. Fernstrom, and Mingui Sun. Imaged based estimation of food volume using circular referents in dietary assessment. *Journal of Food Engineering*, 109(1):76–86, 2012. 1
- [6] Wenyan Jia, Yiqiu Ren, Boyang Li, Britney Beatrice, Jingda Que, Shunxin Cao, Zekun Wu, Zhi-Hong Mao, Benny Lo, Alex K. Anderson, Gary Frost, Megan A. McCrory, Edward Sazonov, Matilda Steiner-Asiedu, Tom Baranowski, Lora E. Burke, and Mingui Sun. A novel approach to dining bowl reconstruction for image-based food volume estimation. *Sensors*, 22(4), 2022.
- [7] Wenyan Jia, Boyang Li, Yaguang Zheng, Zhi-Hong Mao, and Mingui Sun. Estimating amount of food in a circular dining bowl from a single image. In *Proceedings of the 8th International Workshop on Multimedia Assisted Dietary Management*, page 1–9, New York, NY, USA, 2023. Association for Computing Machinery. 1
- [8] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment Anything, 2023. arXiv:2304.02643 [cs]. 1
- [9] Feng Li, Bin He, Gang Li, Zhipeng Wang, and Rong Jiang. Shape-Biased Ellipse Detection Network With Auxiliary Task. *IEEE Transactions on Instrumentation and Measurement*, 71: 1–13, 2022. 1
- [10] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. Grounding dino: Marrying dino with grounded pre-training for open-set object detection, 2023. 2
- [11] Cai Meng, Zhaoxi Li, Xiangzhi Bai, and Fugen Zhou. Arc Adjacency Matrix-Based Fast Ellipse Detection. *IEEE Transactions on Image Processing*, 29:4406–4420, 2020. 1, 4
- [12] Weiqing Min, Shuqiang Jiang, Jitao Sang, Huayang Wang, Xinda Liu, and Luis Herranz. Being a supercook: Joint food attributes and multimodal content modeling for recipe retrieval and exploration. *IEEE Transactions on Multimedia*, 19(5):1100–1113, 2017. 2, 4
- [13] Dilip Prasad. Survey of the problem of object detection in real images. *International Journal of Image Processing (IJIP)*, 6: 441, 2012. 2
- [14] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024. 1
- [15] Stéfán van der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, Tony Yu, and the scikit-image contributors. scikit-image: image processing in python. *PeerJ*, 2: e453, 2014. 4
- [16] Wenyan Jia, Ruizhen Zhao, Ning Yao, John D. Fernstrom, Madelyn H. Fernstrom, Robert J. Scabassi, and Mingui Sun. A food portion size measurement system for image-based dietary assessment. In *2009 IEEE 35th Annual Northeast Bioengineering Conference*, pages 1–2, Cambridge, MA, USA, 2009. IEEE. 1