a



Oregon State University

Garbage In, Knowledge Out

A Smart Compost Bin for Automatic Food Waste Measurement

Qualifying Exam — M.S. Project Defense

Aidan Beery

December 10, 2025

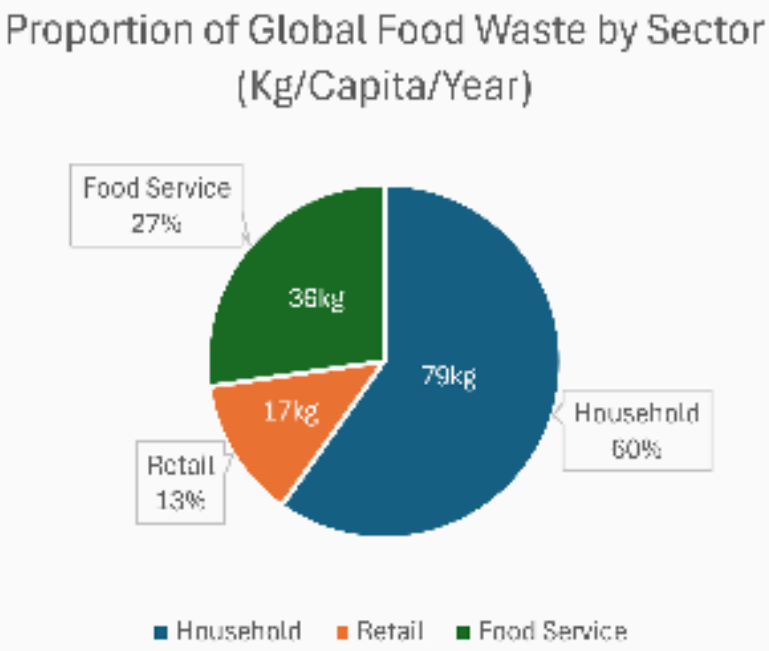
soundbendor lab

Impacts of Food Waste

United Nations SDG 12.3

"By 2030, halve per capita global food waste at the retail and consumer levels and reduce food losses along production and supply chains, including post-harvest losses"

- Over 1 billion tons of food waste & loss annually [6]
- 783 million people food-insecure [7]
- 8-10% of global greenhouse gas emissions stem from FLW
- $1 trillion in economic losses
- **60% of food waste originates from the household** [25]



Global Food Waste, by sector [25]

Food Waste Data Gap

UN FAO Report on SDG Progress

"...Data collection efforts are urgently needed for countries to develop evidence-based, targeted interventions."

- UN FAO 2021: "SDG 12.3 status assessment not possible due to insufficient data" [26]
- Lack of reliable household FW measurements paralyzes sustainability initiatives [31]
- Existing methods are expensive, time-consuming, inaccurate [27]



Hierarchy of food waste disposal behaviors [EPA]

Summary of Contributions

Research Goal

Develop a framework for standardized, scalable, automatic household food waste measurement

1. Design Smart Compost Bin for food waste imaging, measurement, and annotation
2. Begin pilot data collection study with volunteer households
3. Develop novel dataset of annotated food waste images
4. Evaluate CLIP for open-vocabulary food waste image recognition



Smart Compost Bin

Related Works:
Food Waste Measurement

Related Works

Study Type	Procedure	Weaknesses
Mass Balance	Estimate waste from proxy data	Highly inaccurate, insufficient granularity
Question-naires	Self-assessed waste volumes	Users may misrepresent and underreport waste
Waste Composition Analysis	Researchers collect and sort waste	Very high cost, no standard definitions for food categories
Food Waste Diary	Maintain written food waste log	High participant burden, un-derreported waste
Direct Weighing	Weigh waste before disposal	Moderate participant burden, lacks categorical data
Waste Imaging	Take photos of waste before disposal	Requires complex CV models, lacks mass estimation

Comparison of primary food waste measurement approaches.

Related Works: Household Food Waste Surveys

- Manual waste annotation: tedious, burdensome
- Food Waste (FW) Diaries underestimate waste by 7%—40%
- Waste Composition Analysis (WCA) studies don't correlate
  - Lack of standardized waste definitions
  - Lack of standardized measurement protocols
- **Measurement directly reduces wasteful behavior**

Prior Works: Food Imaging

- Imaging for Household FW Measurement
  - Highly correlated to WCA studies [27]
  - Users report reduced annotation burden
  - Leverage smartphones for waste image collection
  - **Current methods require manual image annotation**

Prior Works: Food Imaging

- Imaging for Household FW Measurement
  - Highly correlated to WCA studies [27]
  - Users report reduced annotation burden
  - Leverage smartphones for waste image collection
  - **Current methods require manual image annotation**
- Applications of Automatic Food Image Recognition
  - Nutritional monitoring & health tracking [15]
  - Agricultural automation [18]
  - Food quality control [4]

Prior Works: Commercial Food Waste Imaging

- Measurement reduces institutional FW by 16%—64% [20, 22]
- Leanpath: Image annotation with tablet
- Winnow, Kitro: Uses CV for food waste recognition
  - Closed-source proprietary models
- Base-mounted scale for mass measurement



(a) Winnow [16]



(b) Kitro AI [8]



(c) Leanpath [24]

Examples of imaging-based commercial food waste measurement systems

Prior Works: Smart Bins for Inorganic Waste Measurement

- Smart Bins—IoT waste measurement devices
- Common applications:
  - Organic—inorganic waste separation [23]
  - Automatic separation of recyclables [13]
  - Municipal waste disposal management [21]
  - **Evaluation of waste reduction initiatives**[5]



(a) Consumer Bin [23] **(b)** Recyclables Sorting [13] **(c)** Municipal Bin [21]

Examples of smart waste measurement devices

Designing the Smart Compost Bin

Design

- Key Design Features
  - **Sanitation:** Use removable, dishwasher-safe interior bin
  - **Odor Mitigation:** Passive air filtration with carbon pellets
  - **Convenience:** Automatic speech recognition for disposal annotation
- Manufacturing
  - Final Unit Cost: $468
  - 3D printed design
  - 55 devices produced



Smart Compost Bin (underside)

Design



(a) Baseplate



(b) Body (rear)



(c) Body



(d) Pi Mount



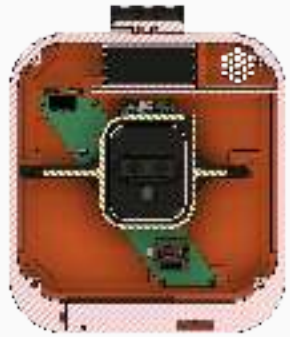
(e) Hinge



(f) Camera



(g) Lid Assembly



(h) Lid Cutaway

3D renders of each component of the smart compost bin

Sensors



(a) Raspberry Pi 5



(b) Depth Camera



(c) Gas Sensor



(d) Amplifier



(e) IR Camera



(f) Mic/DAC



(g) C2 Load Cell

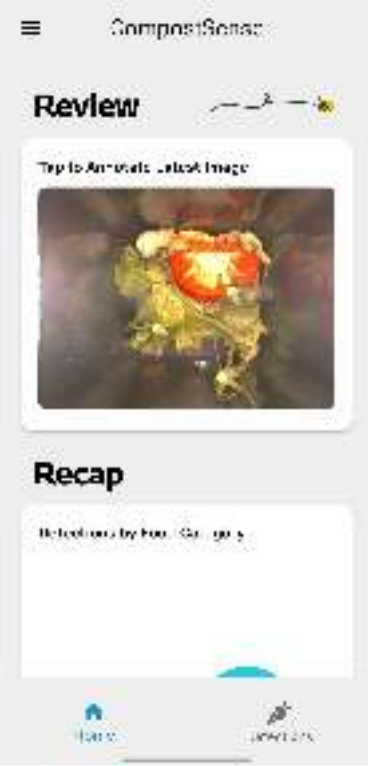
Sensors and microcontroller used in the Smart Compost Bin

CompostSense App

- CompostSense: Mobile app for food waste management
- Pairs with Smart Compost Bin for setup
- Provides analytics dashboard, informing user of waste patterns
- Image annotation feature enables granular FW data collection



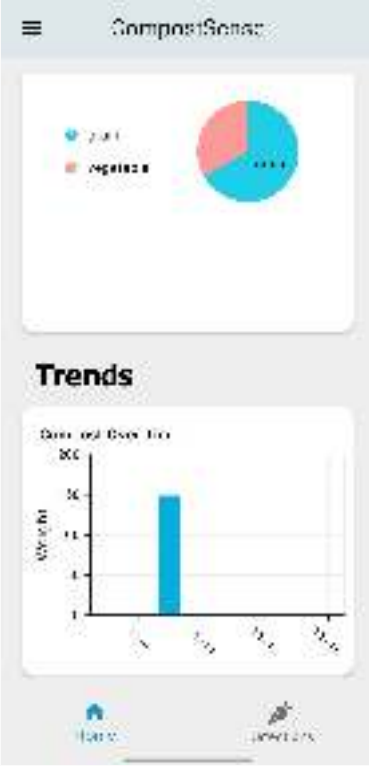
CompostSense App



(a) Home Screen



(b) Waste Groups



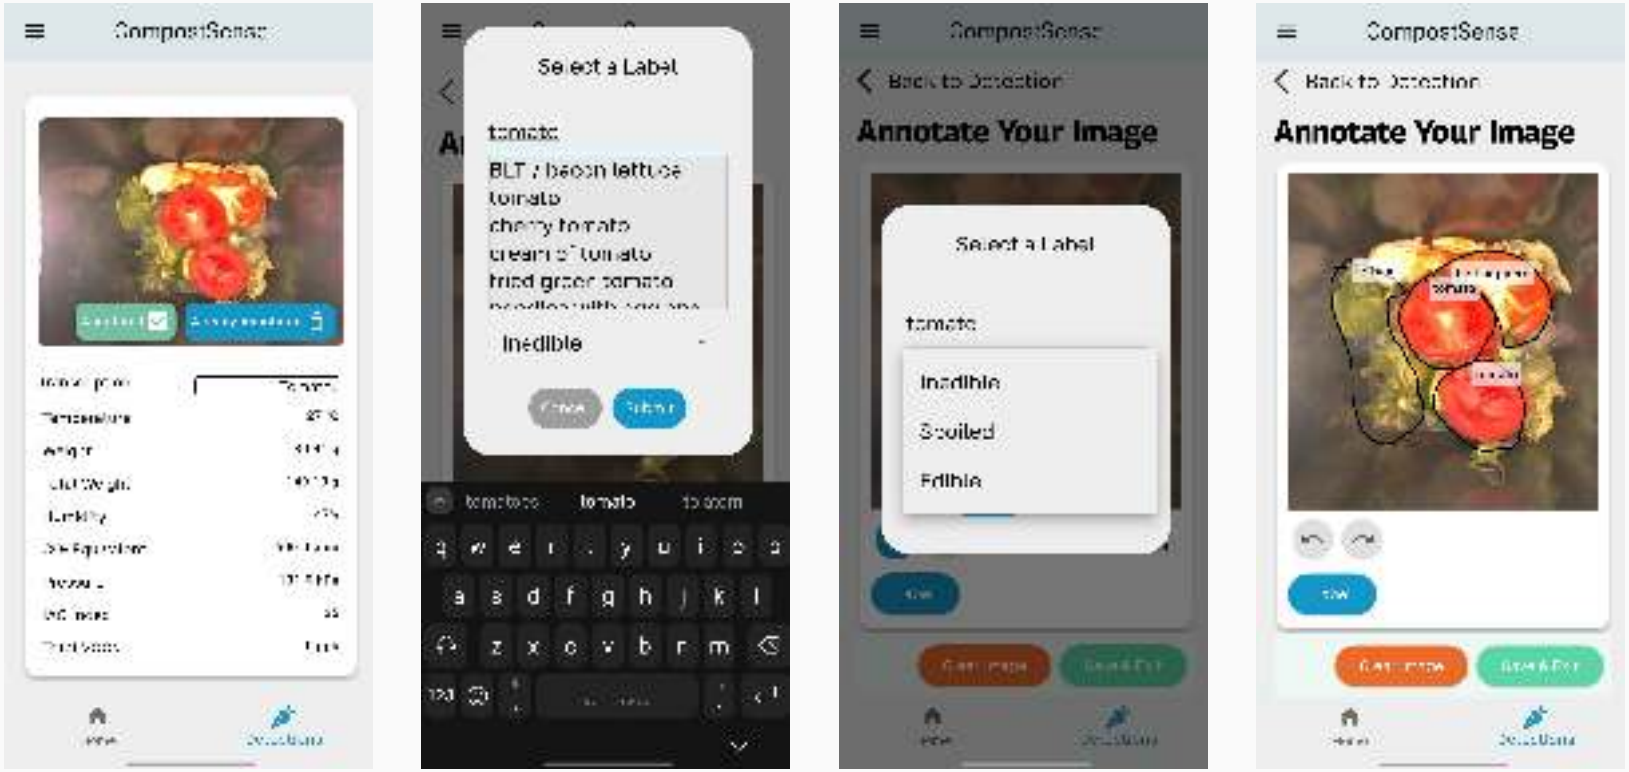
(c) Waste Mass



(d) Detections

User interface for the CompostSense mobile app

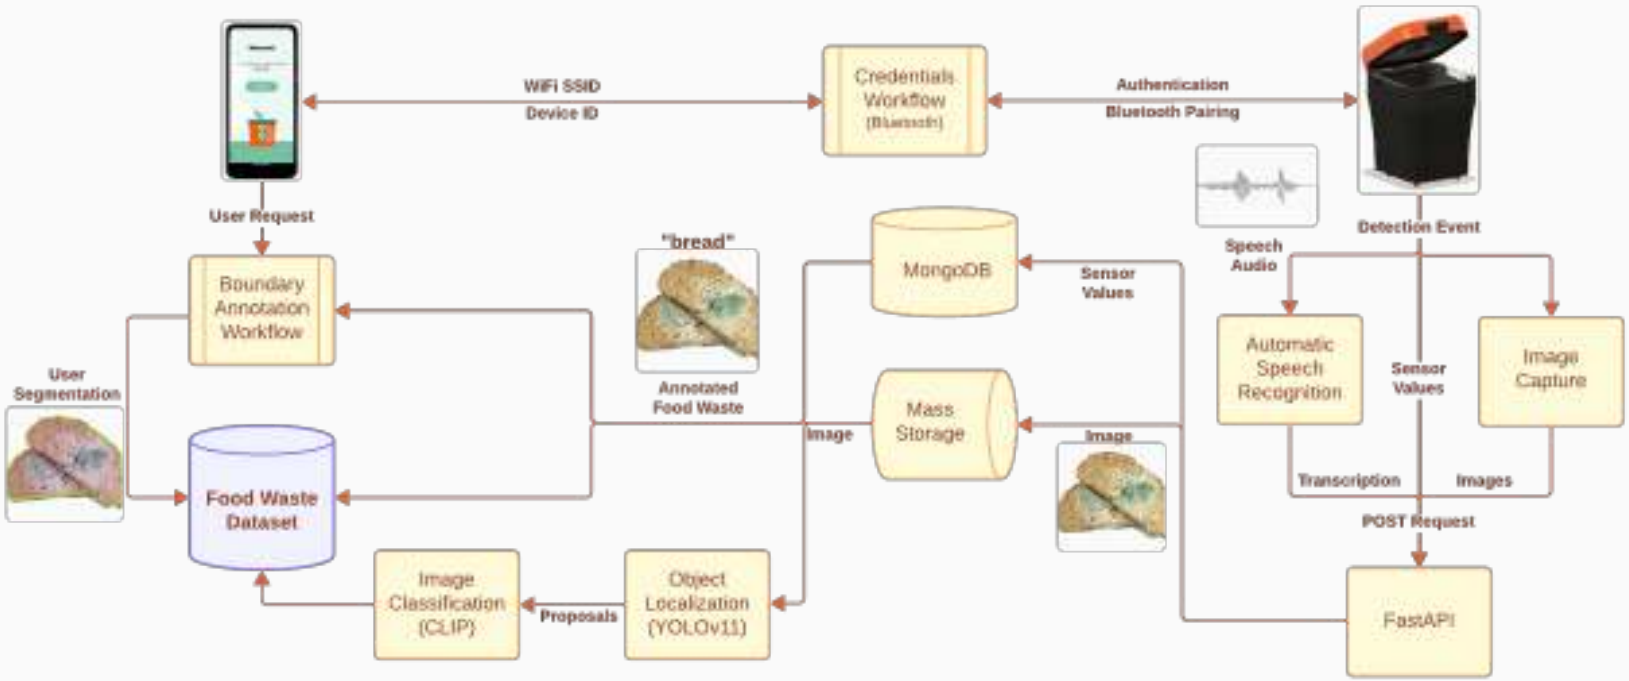
CompostSense Image Annotation



(a) Scan Example (b) Food Label (c) Waste State Tag (d) Labeled Image

Compost image annotation workflow in the CompostSense mobile app

Data Pipeline



System Diagram of the Smart Compost Bin's API server

Food Image Recognition

Summary of Experiments

- Contributions
  1. Evaluate CLIP for compost image recognition
  2. Assess use of food image recognition model on FW tasks
  3. Design framework for open-vocabulary compost object detection

Summary of Experiments

- Contributions
  1. Evaluate CLIP for compost image recognition
  2. Assess use of food image recognition model on FW tasks
  3. Design framework for open-vocabulary compost object detection
- Consider food $\rightarrow$ food waste domain transfer
  - Food image recognition: Analysis of uneaten, plated food.
  - Compost image recognition: Analysis of food waste images

Summary of Experiments

- Contributions
  1. Evaluate CLIP for compost image recognition
  2. Assess use of food image recognition model on FW tasks
  3. Design framework for open-vocabulary compost object detection
- Consider food $\rightarrow$ food waste domain transfer
  - Food image recognition: Analysis of uneaten, plated food.
  - Compost image recognition: Analysis of food waste images
- Consider increasingly difficult image recognition tasks
  - Food Image Classification — 1 : 1 image-label mapping
  - Food Object Classification — Classify each annotated object
  - Food Object Detection — Localize and classify each detection
  - Compost Object Detection — Detect unknown compost items

Challenges in Food Image Recognition

- Food Image Recognition - challenging domain
  - Vast label space
  - Visually similar classes
    - ”Fine-grained” classification [10]
  - Label ambiguity
    - ”Gyoza” v.s. ”Pierogi”
  - Synonymous labels
    - ”Green Onion” v.s. ”Scallion”
  - State Descriptors [3]
    - ”Apple” v.s. ”Apple Slices”

Datasets

- **Problem**: No publicly available dataset of annotated food waste image
- **Goal**: Train food image recognition model, transfer to post-consumer food domain

			Training		Evaluation	
Dataset	**Task**	**# Classes**	# Samples	# Instances	# Samples	# Instances
FoodX-251 [10]	Food Classification	251	118,000	118,000	12,000	12,000
Food201 [15]	Food Object Detection	201	9,653	15,979	2,439	4,149
UECFoodPix [17]	Food Object Detection	102	9,000	13,032	1,000	1,575
FoodSeg103 [29]	Food Object Detection	103	4,983	18,318	2,135	7,697
Combined	Food Object Detection	389	23,637	47,329	5,574	13,421
Food Waste Test	Compost Object Detection	77	–	–	467	1600

Summary of food image datasets, including our new Food Waste Test dataset of annotated images collected using the Smart Compost Bin

CLIP for Food Image Classification

Task: Food Image Classification

Food Image Classification

Labeling of an image based on single relevant food item

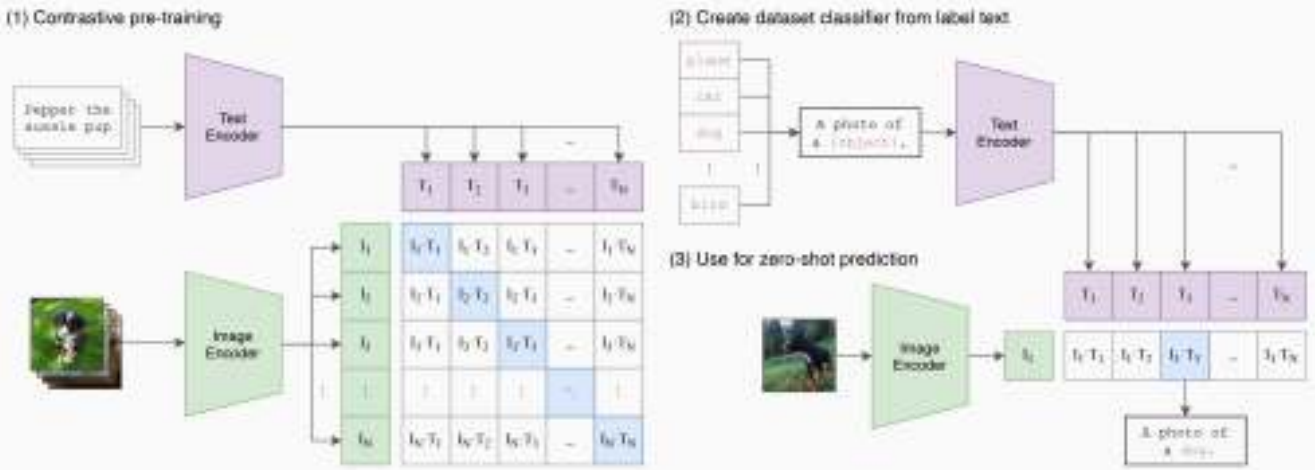
- Datasets
  - Food-101: 101,000 images
  - **FoodX-251:** 130,000 images
- Models
  - CNN (ResNet-101): 0.16 top-3 err. on FoodX-251 [10]
  - ViT (Swin-DR): 81.1% acc. on FoodX-251 [30]
  - CLIP: 95.4% acc. on Food-101 [19]



Example of visually similar classes in the FoodX-251 dataset [10]

Contrastive Image-Language Models

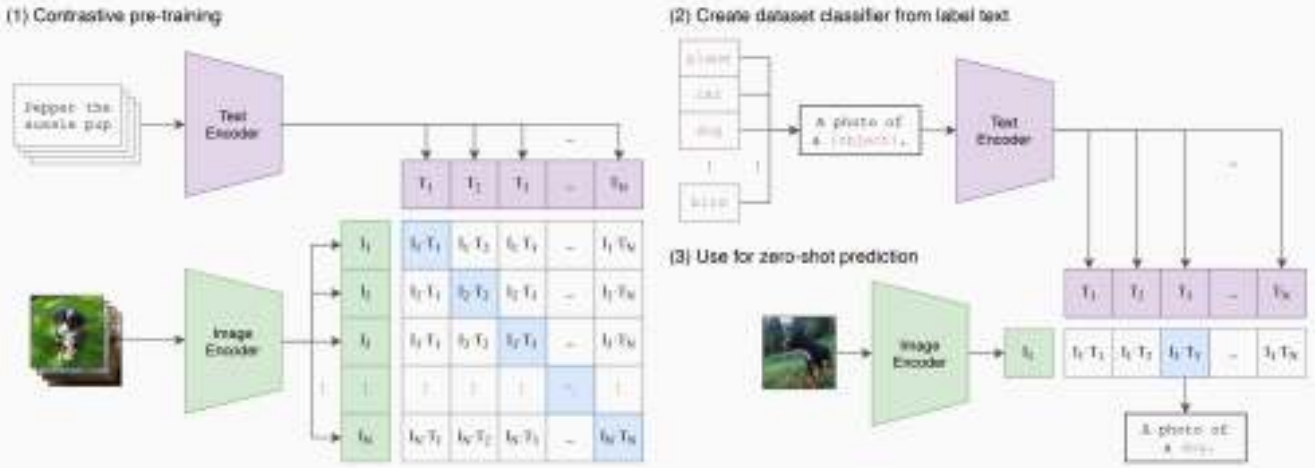
- CLIP: Contrastive Image-Language Pretraining
- Learns joint image-caption feature space
- Generates image embedding vectors
- **Enables open-vocabulary classification**



Contrastive pretraining for multimodal embedding alignment [19]

Contrastive Image-Language Models

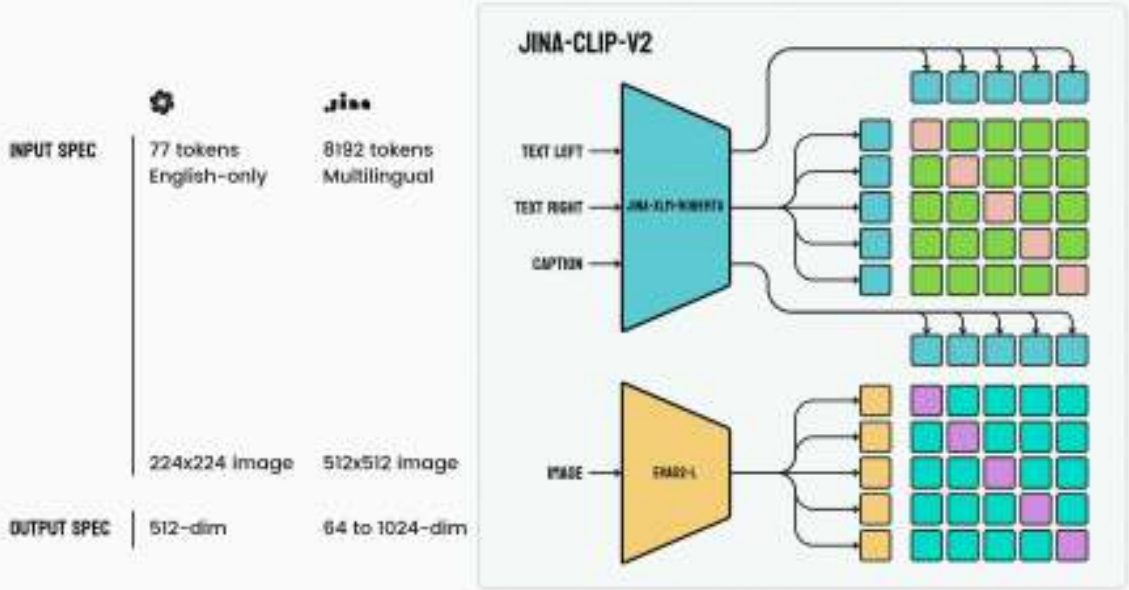
- `CLIP-ViT-B-32`: Reference CLIP implementation [19]
  - 512-dimension embedding vector
  - 32-layer ViT for vision encoder
  - Un-initialized 12-layer text encoder



Contrastive pretraining for multimodal embedding alignment [19]

jina-clip-v2

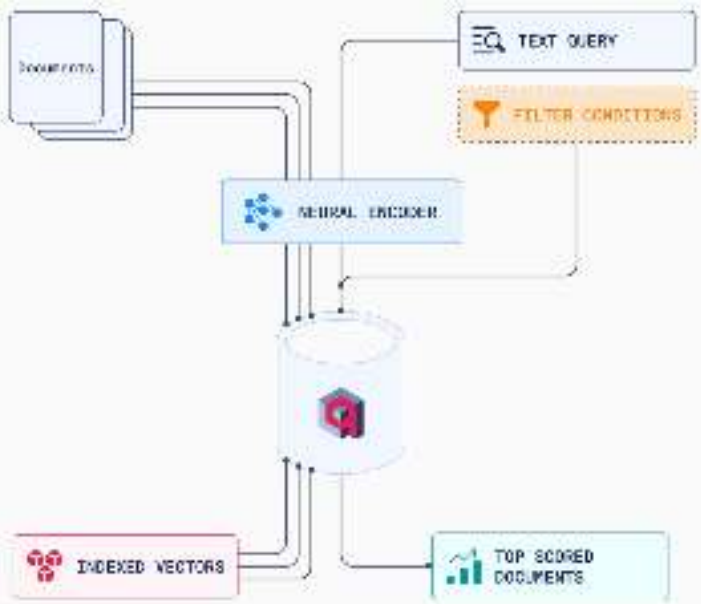
- jina-clip-v2: Open-source CLIP alternative [12]
  - Larger 1024-dimension embeddings
    - Variable-dimension Matryoshka embeddings
  - EVA-02 vision encoder—Robust image classification
  - Pretrained RoBERTa text encoder—natural language grounding



Model architecture of jina-clip-v2

Vector Database: qdrant

- qdrant - Vector Database
- Training: Store each image in database as vector
  - Payload: Class label
- Inference: Query most similar image vectors
  - Vote on predicted class



Workflow for qdrant vector database

Results: Food Image Classification

- Use of larger 1024-dim embeddings yields 17% improvement
- Similar performance to Swin-DR without task-specific finetuning

Model	Output Dim.	Top-1 Acc.	Top-5 Acc	Top-3 Error
CLIP-ViT-B-32	512	0.676	0.838	0.21
jina-clip-v2	512	0.675	0.838	0.21
jina-clip-v2	1024	0.790	0.900	0.15
Swin-DR ViT [30]	—	0.811	—	—
ResNet-101 [10]	—	—	—	0.16

Comparison of CLIP models to current state-of-the-art on FoodX-251.

Results: Food Object Classification

- Goal: Measure transferability between tasks
  - Food image classification ⇒ Food object detection
- Use segmentation labels to compute bounding boxes
- Compute CLIP embeddings for each cropped object

Model	Dataset	Top-1 Acc.	Top-5 Acc.
`CLIP-ViT-B-32` 512-dim. embedding	Food201	0.507	0.651
	UECFoodPix	0.686	0.864
	FoodSeg103	0.500	0.664
	Combined	0.466	0.632
`jina-clip-v2` 1024-dim. embedding	Food201	0.584	0.724
	UECFoodPix	0.741	0.864
	FoodSeg103	0.581	0.738
	Combined	0.521	0.683

Evaluation of CLIP models for food object classification.

Results: Food Object Classification

- Goal: Measure transferability between tasks
  - Food image classification ⇒ Food object detection
- Use segmentation labels to compute bounding boxes
- Compute CLIP embeddings for each cropped object
- Result: ≈ 23% decrease in accuracy

Model	Dataset	Top-1 Acc.	Top-5 Acc.
`CLIP-ViT-B-32` 512-dim. embedding	Food201	0.507	0.651
	UECFoodPix	0.686	0.864
	FoodSeg103	0.500	0.664
	Combined	0.466	0.632
`jina-clip-v2` 1024-dim. embedding	Food201	0.584	0.724
	UECFoodPix	0.741	0.864
	FoodSeg103	0.581	0.738
	Combined	0.521	0.683

Evaluation of CLIP models for food object classification.

Food Object Detection

Task: Food Object Detection

Food Object Detection

Localization and classification of multiple food items in an image

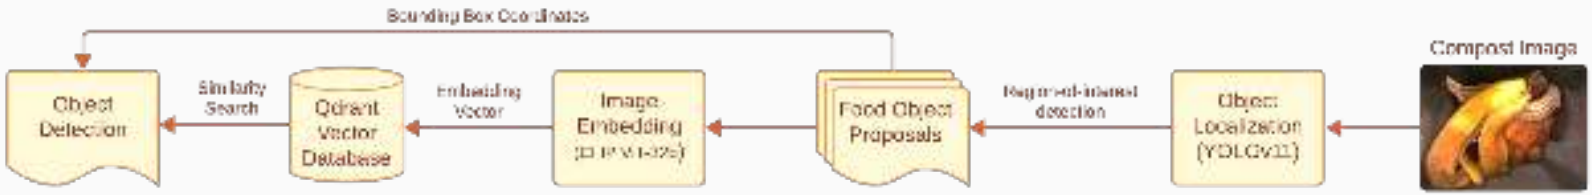
- YOLO: CNN for single-pass object detection
  - Broadly applied for food object detection tasks
  - Fine-tuning requires relatively few samples
- Strong performance in downstream food object detection tasks [9, 2]
- **Limited by fixed label space**



Example of converted bounding-boxes from UECFoodPix [17] and FoodSeg103 [29]

Food Object Detection Pipeline

- **Training Pipeline**
  - Fine-tune YOLOv11 on food object detection datasets
  - Generate embeddings from annotator-cropped objects
  - Store embeddings in vector database
- **Inference**
  - Use YOLOv11 to detect regions-of-interest
  - Generate CLIP embeddings from proposals
  - Use vector similarity for proposal classification



Two-stage inference pipeline for open-vocabulary food object detection

Results: YOLOv11 Baseline

- Fine-tune YOLOv11 for object localization
- Use classification results as baseline
- Fine-tune for 85 epochs, mAP early stopping criteria

Dataset	Precision	Recall	mAP50	mAP50-95
Food201	0.524	0.467	0.340	0.284
UECFoodPix	0.692	0.687	0.200	0.186
FoodSeg103	0.589	0.554	0.342	0.296
Combined	0.584	0.545	0.293	0.251

Performance of a fine-tuned YOLOv11 model on food object detection tasks.

Results: YOLOv11 + CLIP Food Object Detection

- Use YOLOv11 to detect proposals
- Embeddings from `jina-clip-v2`

Model	Dataset	Precision	Recall	mAP50	mAP50-95
`jina-clip-v2`	Food201	0.446	0.397	0.297	0.261
	UECFoodPix	0.652	0.648	0.188	0.178
	FoodSeg103	0.534	0.503	0.304	0.263
	Combined	0.456	0.425	0.226	0.197

Evaluation of YOLOv11 + `jina-clip-v2` for food object detection.

Results: YOLOv11 + CLIP Food Object Detection

- YOLO + CLIP does not improve performance for fixed label space
- 13.61% decrease in average mAP50
  - YOLOv11: 0.294
  - YOLO + CLIP: 0.254

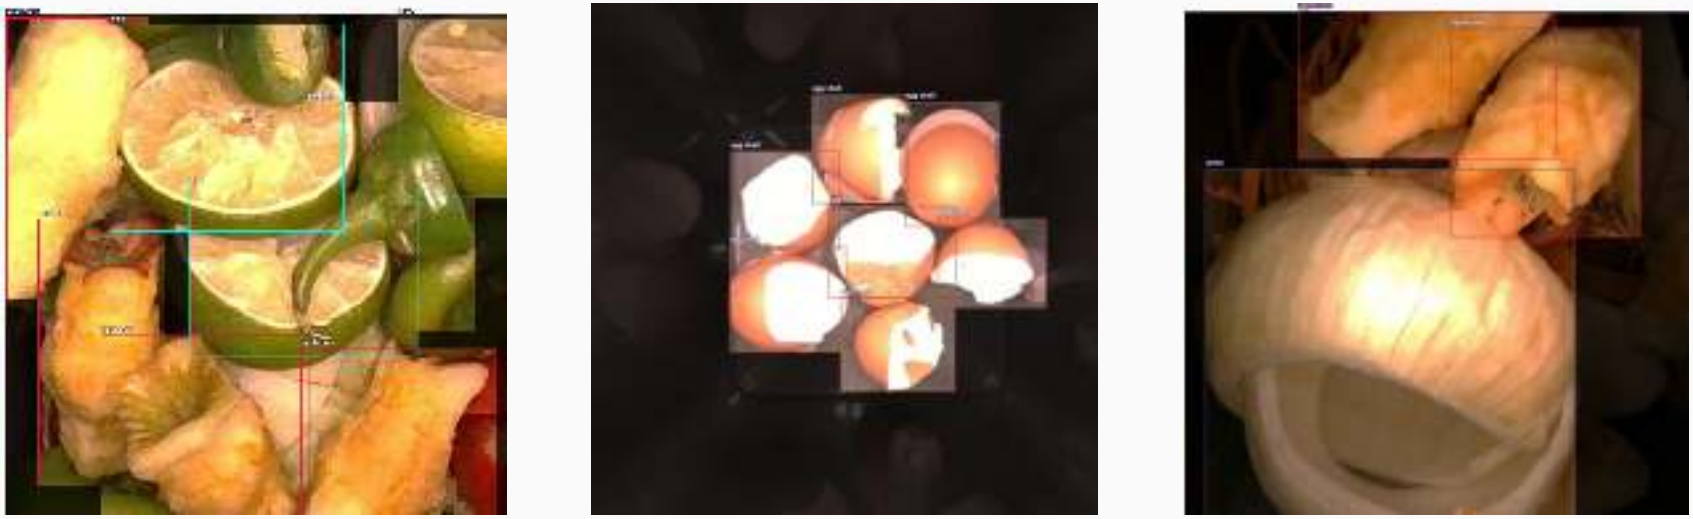
Model	Dataset	Precision	Recall	mAP50	mAP50-95
`jina-clip-v2`	Food201	0.446	0.397	0.297	0.261
	UECFoodPix	0.652	0.648	0.188	0.178
	FoodSeg103	0.534	0.503	0.304	0.263
	Combined	0.456	0.425	0.226	0.197

Evaluation of YOLOv11 + `jina-clip-v2` for food object detection.

Compost Object Detection

Task: Compost Object Detection

- Evaluate model on open-vocabulary compost object detection
- Use food waste images captured from Smart Compost Bin
  - 467 images, 1600 instances
  - 77 classes, preserved food state descriptors
- Add text embeddings for labels into vector database



Examples of annotated images from the Smart Compost Bin

Results: Compost Object Detection

- Achieve modest precision of 0.291
- Maximum theoretical recall: 0.18
  - YOLOv11 fails to detect most objects in scene
- Fine-tuning on food image datasets does not immediately yield a stronger compost object detection model
- Demonstrates feasibility of using CLIP for open-vocabulary compost object detection tasks

Model	Precision	Recall	mAP50	mAP50-95
`jina-clip-v2`	0.291	0.116	0.063	0.076
`jina-clip-v2-foodx251`	0.266	0.106	0.059	0.052
`jina-clip-v2-combined`	0.180	0.072	0.048	0.032

Evaluation of YOLOv11 + `jina-clip-v2` for compost object detection.

Discussion

Discussion - Smart Compost Bin

- Smart Compost Bin streamlines food waste measurement
  - Unify weighing and imaging steps
  - Allow users to dictate waste labels at disposal time
  - Eliminate need for food waste diaries

Discussion - Smart Compost Bin

- Smart Compost Bin streamlines food waste measurement
  - Unify weighing and imaging steps
  - Allow users to dictate waste labels at disposal time
  - Eliminate need for food waste diaries
- Integrated software ecosystem facilitates waste measurement
  - CompostSense app provides user interface for compost bin
    - Enables detailed instance-level segmentation annotations
    - Food waste analytics dashboard helps users reduce waste
  - Database & API server centralize data collection and analysis

Discussion - Smart Compost Bin

- Smart Compost Bin streamlines food waste measurement
  - Unify weighing and imaging steps
  - Allow users to dictate waste labels at disposal time
  - Eliminate need for food waste diaries
- Integrated software ecosystem facilitates waste measurement
  - CompostSense app provides user interface for compost bin
    - Enables detailed instance-level segmentation annotations
    - Food waste analytics dashboard helps users reduce waste
  - Database & API server centralize data collection and analysis
- Scalable device deployment enables standardized household FW measurement studies
  - Imaging enables waste composition analysis
  - Weight sensor directly measures waste volumes

Discussion - Food Waste Image Recognition

- CLIP shows potential for compost image recognition
  - CLIP attains near-SoTA performance on FoodX-251 dataset
  - 23% performance penalty when transferred to object classification
  - Some capability on previously unseen compost image data

Discussion - Food Waste Image Recognition

- CLIP shows potential for compost image recognition
  - CLIP attains near-SoTA performance on FoodX-251 dataset
  - 23% performance penalty when transferred to object classification
  - Some capability on previously unseen compost image data
- YOLO + CLIP not suitable for food object recognition tasks
  - YOLO + CLIP does not exceed YOLO alone for food object detection
  - Low theoretical recall—YOLO fails to detect food waste items
  - Tradeoff: 13.61% mAP50 decrease for open-vocabulary capability

Discussion - Food Waste Image Recognition

- CLIP shows potential for compost image recognition
  - CLIP attains near-SoTA performance on FoodX-251 dataset
  - 23% performance penalty when transferred to object classification
  - Some capability on previously unseen compost image data
- YOLO + CLIP not suitable for food object recognition tasks
  - YOLO + CLIP does not exceed YOLO alone for food object detection
  - Low theoretical recall—YOLO fails to detect food waste items
  - Tradeoff: 13.61% mAP50 decrease for open-vocabulary capability
- Food image fine-tuning insufficient for compost object detection
  - Base `jina-clip-v2` best on food waste evaluation
  - Cropped food objects may be incompatible with CLIP fine-tuning

Limitations

- Food Waste Test dataset limited in size
  - 6 bins deployed in beta testing
  - Full-scale data collection study forthcoming
- YOLOv11 constrains CLIP performance
- Bounding boxes lack granularity for food image recognition
  - Compost images highly commingled, noisy boxes
- CLIP lacks spatial capabilities
  - Pre-trained on image captioning tasks
  - Can not perform localization natively
  - Does not fine-tune well with food objects

Future Task: Food Instance Segmentation

Food Instance Segmentation

Pixel-level localization and classification of food items in an image

- Segmentations provide specificity in visually complex domains
- Pretrained foundation segmentation models [11]
- Fine-tuning SAM for downstream tasks still an open problem [32]
- **Goal**: Use CLIP for open-vocabulary compost instance segmentation



Example of segmentation annotations from Food-201 [15]

Future Work

- Embed spatial capabilities into CLIP
  - Retain open-vocabulary classification capability
  - SAM-CLIP: Knowledge distillation to merge SAM into CLIP [28]
  - Enables end-to-end open-vocabulary instance segmentation

Future Work

- Embed spatial capabilities into CLIP
  - Retain open-vocabulary classification capability
  - SAM-CLIP: Knowledge distillation to merge SAM into CLIP [28]
  - Enables end-to-end open-vocabulary instance segmentation
- Experiment with VLMs for compost image recognition
  - VLMs still struggle with image classification tasks [33]
  - Emerging work in instruction-tuning for spatial tasks [1]
  - Instruction tuning may enable compost recognition

Future Work

- Embed spatial capabilities into CLIP
  - Retain open-vocabulary classification capability
  - SAM-CLIP: Knowledge distillation to merge SAM into CLIP [28]
  - Enables end-to-end open-vocabulary instance segmentation
- Experiment with VLMs for compost image recognition
  - VLMs still struggle with image classification tasks [33]
  - Emerging work in instruction-tuning for spatial tasks [1]
  - Instruction tuning may enable compost recognition
- Augment image similarity classification with reranker models

Future Work

- Embed spatial capabilities into CLIP
  - Retain open-vocabulary classification capability
  - SAM-CLIP: Knowledge distillation to merge SAM into CLIP [28]
  - Enables end-to-end open-vocabulary instance segmentation
- Experiment with VLMs for compost image recognition
  - VLMs still struggle with image classification tasks [33]
  - Emerging work in instruction-tuning for spatial tasks [1]
  - Instruction tuning may enable compost recognition
- Augment image similarity classification with reranker models
- Extension: Recipe retrieval with Recipe1M+ [14]
  - Embed recipes as documents in vector database
  - Retrieve recipes based on most frequently wasted items
  - Nudge users towards more sustainable consumption

Thank you!

We are recruiting participants for this study. If you are interested in volunteering to use the Smart Compost Bin, please scan this QR code to fill out our interest survey.

References

[1] Lucas Beyer et al. ***PaliGemma: A versatile 3B VLM for transfer.*** 2024. arXiv: 2407.07726 [cs.CV]. URL: https://arxiv.org/abs/2407.07726.

[2] Gustav Chrintz-Gath et al. **"Image-Based Dietary Assessment Using the Swedish Plate Model: Evaluation of Deep Learning-Based You Only Look Once (YOLO) Models".** In: *JMIR formative research* 9 (Aug. 14, 2025), e70124. ISSN: 2561-326X. DOI: 10.2196/70124. PMID: 40812290.

[3] Gianluigi Ciocca, Giovanni Micali, and Paolo Napoletano. **"State Recognition of Food Images Using Deep Features".** In: *IEEE Access* 8 (2020), pp. 32003–32017. DOI: 10.1109/ACCESS.2020.2973704.

References ii

[4] Ayushi Dhelia, Srishti Chordia, and Kanisha B. **"YOLO-based Food Damage Detection: An Automated Approach for Quality Control in Food Industry".** In: *2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*. 2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC). Oct. 2024, pp. 1444–1449. DOI: 10.1109/I-SMAC61858.2024.10714664. URL: https://ieeexplore.ieee.org/abstract/document/10714664 (visited on 09/26/2025).

[5] Bingbing Fang et al. **"Artificial Intelligence for Waste Management in Smart Cities: A Review".** In: *Environmental Chemistry Letters* (May 9, 2023), pp. 1–31. ISSN: 1610-3653. DOI: 10.1007/s10311-023-01604-3. PMID: 37362015. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10169138/ (visited on 03/19/2024).

References iii

[6] Food and Agriculture Organization (FAO). **"Food wastage footprint: Impacts on natural resources".** In: *Summary Report* (2013), pp. 1–63.

[7] Hamish Forbes, Tom Quested, and Clemintine O'Connor. **"Food Waste Index Report".** In: *United Nations Environment Programme* (2021).

[8] Giacinto Bottone. ***Start-up of the Day: Battling Food Waste with AI.*** IO+ Archive. Apr. 8, 2020. URL: https://ioplus.nl/archive/en/start-up-of-the-day-battling-food-waste-with-ia/ (visited on 12/07/2025).

[9] Aknur Karabay, Huseyin Atakan Varol, and Mei Yen Chan. **"Improved Food Image Recognition by Leveraging Deep Learning and Data-Driven Methods with an Application to Central Asian Food Scene".** In: *Scientific Reports* 15.1 (Apr. 23, 2025), p. 14043. ISSN: 2045-2322. DOI: 10.1038/s41598-025-95770-9. URL: https://www.nature.com/articles/s41598-025-95770-9 (visited on 09/26/2025).

References iv

[10] Parneet Kaur et al. ***FoodX-251: A Dataset for Fine-grained Food Classification.*** July 14, 2019. DOI: 10.48550/arXiv.1907.06167. arXiv: 1907.06167 [cs]. URL: http://arxiv.org/abs/1907.06167 (visited on 03/31/2025). Pre-published.

[11] Alexander Kirillov et al. ***Segment Anything.*** 2023. arXiv: 2304.02643 [cs.CV]. URL: https://arxiv.org/abs/2304.02643.

[12] Andreas Koukounas et al. ***jina-clip-v2: Multilingual Multimodal Embeddings for Text and Images.*** 2025. arXiv: 2412.08802 [cs.CL]. URL: https://arxiv.org/abs/2412.08802.

[13] Edoardo Longo et al. **“Take the Trash out... to the Edge. Creating a Smart Waste Bin Based on 5G Multi-access Edge Computing”.** In: *Proceedings of the Conference on Information Technology for Social Good*. GoodIT ’21. New York, NY, USA: Association for Computing Machinery, Sept. 9, 2021, pp. 55–60. ISBN: 978-1-4503-8478-0. DOI: 10.1145/3462203.3475921. URL: https://dl.acm.org/doi/10.1145/3462203.3475921 (visited on 04/04/2025).

[14] Javier Marin et al. ***Recipe1M+: A Dataset for Learning Cross-Modal Embeddings for Cooking Recipes and Food Images.*** 2019. arXiv: 1810.06553 [cs.CV]. URL: https://arxiv.org/abs/1810.06553.

[15] Austin Myers et al. **“Im2Calories: Towards an Automated Mobile Vision Food Diary”.** In: *2015 IEEE International Conference on Computer Vision (ICCV)*. 2015 IEEE International Conference on Computer Vision (ICCV). Dec. 2015, pp. 1233–1241. DOI: 10.1109/ICCV.2015.146. URL: https://ieeexplore.ieee.org/document/7410503 (visited on 03/31/2025).

[16] Yu Nu, Elena Belavina, and Karan Girotra. **“Using artificial intelligence to reduce food waste”.** In: *Available at SSRN 4826777* (2024).

[17] Kaimu Okamoto and Keiji Yanai. **"UEC-FoodPix Complete: A Large-Scale Food Image Segmentation Dataset".** In: *Pattern Recognition. ICPR International Workshops and Challenges*. Ed. by Alberto Del Bimbo et al. Vol. 12665. Cham: Springer International Publishing, 2021, pp. 647–659. ISBN: 978-3-030-68820-2 978-3-030-68821-9. DOI: 10.1007/978-3-030-68821-9_51. URL: http://link.springer.com/10.1007/978-3-030-68821-9_51 (visited on 03/31/2025).

[18] Diego Inácio Patrício and Rafael Rieder. **"Computer Vision and Artificial Intelligence in Precision Agriculture for Grain Crops: A Systematic Review".** In: *Computers and Electronics in Agriculture* 153 (2018), pp. 69–81. ISSN: 0168-1699. DOI: 10.1016/j.compag.2018.08.001. URL: https://www.sciencedirect.com/science/article/pii/S0168169918305829.

References viii

[19] Alec Radford et al. ***Learning Transferable Visual Models From Natural Language Supervision.*** 2021. arXiv: 2103.00020 [cs.CV]. URL: https://arxiv.org/abs/2103.00020.

[20] Steve Rock. **"Food Waste Reduction in Military Kitchens, A Tracking Technology Demonstration at Fort Jackson".** In: (Sept. 2019).

[21] Alberto Rovetta et al. **"Early Detection and Evaluation of Waste through Sensorized Containers for a Collection Monitoring Application".** In: *Waste Management* 29.12 (Dec. 1, 2009), pp. 2939–2949. ISSN: 0956-053X. DOI: 10.1016/j.wasman.2009.08.016. URL: https://www.sciencedirect.com/science/article/pii/S0956053X09003511 (visited on 04/04/2025).

[22] Kirsi Silvennoinen et al. **“Food Waste Volume and Origin: Case Studies in the Finnish Food Service Sector”.** In: *Waste Management (New York, N.Y.)* 46 (Dec. 2015), pp. 140–145. ISSN: 1879-2456. DOI: 10.1016/j.wasman.2015.09.010. PMID: 26419775.

[23] Florian Taurer et al. **“Smart Trash Can: Easy Collection of Photos of Organic Waste in the Home”.** In: *Proceedings of the 22nd International Conference on Mobile and Ubiquitous Multimedia.* MUM ’23: International Conference on Mobile and Ubiquitous Multimedia. Vienna Austria: ACM, Dec. 3, 2023, pp. 571–573. ISBN: 979-8-4007-0921-0. DOI: 10.1145/3626705.3631881. URL: https://dl.acm.org/doi/10.1145/3626705.3631881 (visited on 04/04/2025).

[24] Tekla Perry. ***Leanpath Weighs Waste - IEEE Spectrum.*** May 31, 2013. URL: https://spectrum.ieee.org/leanpath-weighs-waste (visited on 12/07/2025).

[25] United Nations Environment Programme. ***Food Waste Index Report 2024. Think Eat Save: Tracking Progress to Halve Global Food Waste.*** Mar. 2024. ISBN: 978-92-807-4139-1. URL: https://wedocs.unep.org/xmlui/handle/20.500.11822/45230 (visited on 01/16/2025).

[26] United Nations Food and Agriculture Organization. ***Tracking Progress on Food and Agriculture-Related SDG Indicators 2021: A Report on the Indicators under FAO Custodianship.*** Sept. 22, 2021. ISBN: 978-92-5-134967-0. DOI: 10.4060/cb6872en. URL: http://www.fao.org/documents/card/en/c/CB6872EN (visited on 12/15/2024).

[27] Erica van Herpen et al. **"Comparing Wasted Apples and Oranges: An Assessment of Methods to Measure Household Food Waste".** In: *Waste Management* 88 (Apr. 1, 2019), pp. 71–84. ISSN: 0956-053X. DOI: 10.1016/j.wasman.2019.03.013. URL: https://www.sciencedirect.com/science/article/pii/S0956053X19301345 (visited on 01/27/2025).

[28] Haoxiang Wang et al. ***SAM-CLIP: Merging Vision Foundation Models towards Semantic and Spatial Understanding.*** 2024. arXiv: 2310.15308 [cs.CV]. URL: https://arxiv.org/abs/2310.15308.

[29] Xiongwei Wu et al. ***A Large-Scale Benchmark for Food Image Segmentation.*** May 12, 2021. DOI: 10.48550/arXiv.2105.05409. arXiv: 2105.05409 [cs]. URL: http://arxiv.org/abs/2105.05409 (visited on 03/31/2025). Pre-published.

[30] Zhiyong Xiao, Guang Diao, and Zhaohong Deng. **"Fine Grained Food Image Recognition Based on Swin Transformer".** In: *Journal of Food Engineering* 380 (2024), p. 112134. ISSN: 0260-8774. DOI: 10.1016/j.jfoodeng.2024.112134. URL: https://www.sciencedirect.com/science/article/pii/S0260877424002000.

[31] Li Xue et al. **"Missing Food, Missing Data? A Critical Review of Global Food Losses and Food Waste Data".** In: *Environmental Science & Technology* 51.12 (May 2017), pp. 6618–6633. ISSN: 0013-936X. DOI: 10.1021/acs.est.7b00401. (Visited on 04/09/2024).

[32] Yang Yu, Chen Xu, and Kai Wang. ***TS-SAM: Fine-Tuning Segment-Anything Model for Downstream Tasks.*** 2024. arXiv: 2408.01835 [cs.CV]. URL: https://arxiv.org/abs/2408.01835.

[33] Yifan Zhang et al. ***Deep Long-Tailed Learning: A Survey.*** 2023. arXiv: 2110.04596 [cs.CV]. URL: https://arxiv.org/abs/2110.04596.