Italian Journal of Pediatric Allergy and Immunology number 02 | 2025 | 30-33

SIAIP

An intelligent analysis of food allergens through computer vision and generative models

Original article

Alessandra Perrotta[1], Gianluca Mondillo[1], Mariapia Masino[1], Simone Colosimo[1], Vittoria Frattolillo[1], Agnese Sara Ciccarelli[1], Cristiana Indolfi[1], Michele Miraglia del Giudice[1]

[1] Department of Woman, Child and of General and Specialized Surgery, Università degli Studi della Campania "Luigi Vanvitelli", Naples, Italy

SUMMARY

Introduction. Food allergies represent a leading cause of adverse reactions and hospital admissions among children, with significant impact on quality of life and public health. The rapid and accurate detection of allergens in meals is therefore crucial for safety.

Materials and methods. We developed an AI-based prototype that combines YOLOv8n, a state-of-the-art object detection model trained on the Allergen30 dataset, with Gemini 2.0 Flash, an advanced generative model, to provide multimodal allergen analysis. All images were preprocessed and split into training (70%), validation (15%), and test (15%) sets, with careful class balancing.

Results. The system achieved high class-specific performance in detecting allergenic foods from real meal images, with mAP50 >90% and detailed contextual analysis via Gemini 2.0 Flash.

Discussion and conclusions. AI-assisted allergen analysis from meal images is feasible and shows promise, but does not replace ingredient disclosure or clinical precautions. Further development and real-world validation are warranted.

KEYWORDS: Food allergy, Artificial intelligence, YOLOv8n, Gemini 2.0 Flash, Computer vision

Received: May 26, 2025
Published: July 28, 2025

CORRESPONDENCE

Alessandra Perrotta
alessandraperrotta96@gmail.com

How to cite this article: Perrotta A, Mondillo G, Masino M, et al. An intelligent analysis of food allergens through computer vision and generative models. Italian Journal of Pediatric Allergy and Immunology 2025;39(02):30-33. https://doi.org/10.53151/2531-3916/2025-1375

INTRODUCTION

Food allergies affect up to 8% of children worldwide, accounting for a major share of food-related anaphylactic events and adverse reactions in the pediatric population [1,10]. The early and precise identification of potential allergens in daily meals is critical for both individual and public safety, especially for children with severe or multiple allergies [2,11]. Traditional approaches, such as careful label reading, manual inspection, or reliance on ingredient lists, are frequently insufficient due to human error, incomplete information, and the presence of hidden or cross- contact allergens [12,13]. The recent advancement of artificial intelligence (AI) and computer vision offers new solutions for automated food recognition and safety assessment [3,14,15]. Deep learning models such as YOLOv8 have demonstrated remarkable accuracy in object detection and classifying food items from images [16]. However, visual recognition alone cannot identify invisible or trace allergens, nor can it account for complex preparation methods, contamination, or ingredient substitutions [17]. To overcome these limitations, we present a novel prototype that integrates YOLOv8 for object detection with Gemini 2.0 Flash, a cutting-edge generative AI model capable of contextual inference from multimodal input [5,6,18]. This hybrid approach aims to provide not only robust visual classification but also inferential warnings about hidden or probable allergens, based on both detected items and external knowledge [7,19]. The objective of this study is to describe the technical development and preliminary evaluation of this AI- powered allergen analysis

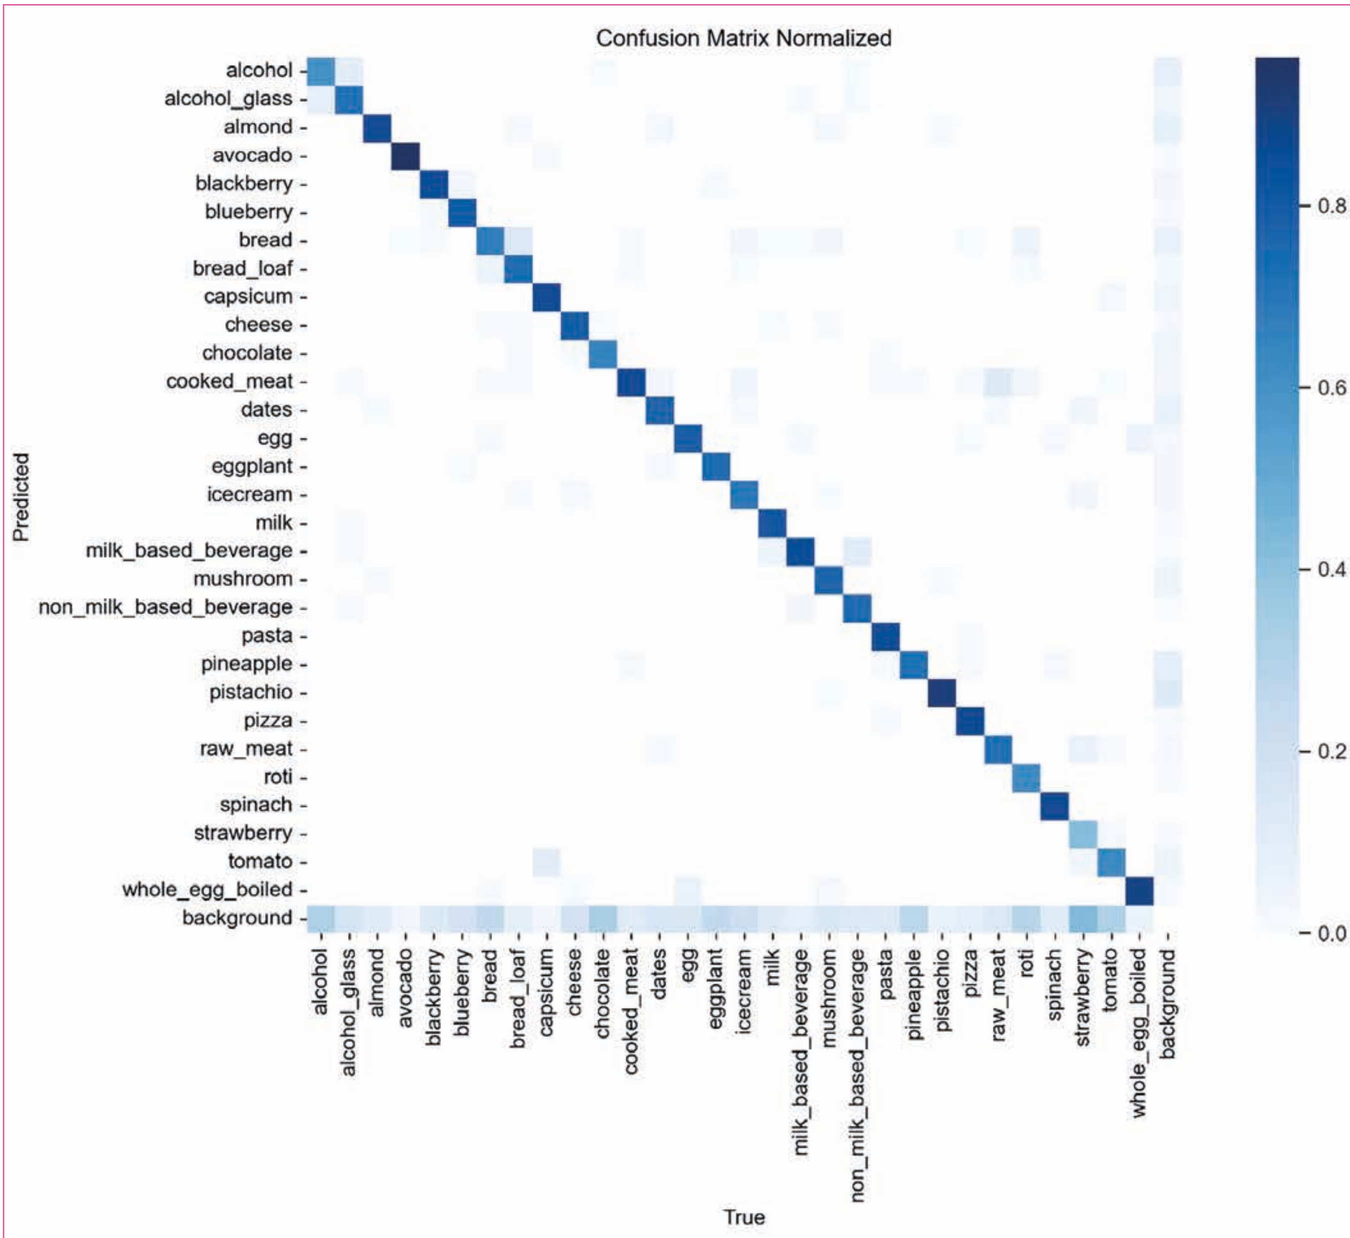

FIGURE 1. Normalized confusion matrix showing class-specific detection performance of the YOLOv8n model on the Allergen30 test set. The matrix highlights high accuracy and low rates of misclassification for the majority of allergenic food classes analyzed.

system, focusing on its potential utility for food-allergic children and their caregivers in home and public settings [8].

MATERIALS AND METHODS

The core dataset used for allergen recognition was Allergen30, a curated collection of 3,000 annotated images encompassing 30 common allergenic food categories including nuts, shellfish, dairy, eggs, and wheat. These categories were selected based on prevalence data in pediatric populations [3]. Each image was manually annotated by trained reviewers to identify allergenic ingredients, with cross-checking to ensure accuracy. All images were resized to 416x416 pixels and converted to the YOLO label format. Dataset splitting into training (70%), validation (15%), and test (15%) sets was performed with a custom algorithm to ensure at least 10 images per class for training and a minimum of 3 per class for validation and test. Manual review and cleaning of labels and images were conducted to maximize data integrity [20].

The YOLOv8n model was trained with the following hyperparameters: 10 epochs, batch size 32, image size 416x416 px, AdamW optimizer, initial learning rate 0.001, final learning rate 0.01, weight decay 0.0005, 3 warmup epochs, and no dropout. The training was performed on a workstation equipped with an Intel Core i9-11900KF 3.50 GHz CPU, NVIDIA RTX 4060 Ti GPU (16 GB VRAM), 32 GB RAM, and SSD storage running Windows 11. Model performance was evaluated using precision, recall, F1-score, and mAP50.

After object detection, Gemini 2.0 Flash (Google DeepMind) was used through API calls to provide contextual, generative analysis for each detected item. Gemini employs a multimodal transformer architecture to analyze detected food items in combination with metadata and prior knowledge. For each image, Gemini generates a risk report flagging likely allergens, cross-contamination risks, and preparation-related hazards. This report is grounded in the co-occurrence of food classes, known recipes, and typical allergen presence associated with the identified items [6,22].

The user interface was built with Gradio for intuitive photo upload, instant image annotation, and generation of a detailed textual risk report [23].

RESULTS

Evaluation of the prototype using the test set from Allergen30 demonstrated high detection accuracy across most allergen classes, with precision and recall typically above 0.85 [3,4,16]. The normalized confusion matrix (Fig. 1) revealed that the majority of classes had very low misclassification rates [4,16]. Mean average precision (mAP50) for the system exceeded 90% on the test set [4,16,21]. Visual inspection of model outputs confirmed reliable localization and identification of multiple allergenic items within composite meal images [21]. Gemini 2.0 Flash provided a contextual narrative for each image, flagging hidden risks such as likely nut traces in desserts or probable milk in baked goods [6,22,24]. The Gradio interface enabled rapid and user-friendly image upload and report generation, with no significant usability issues in initial user tests [23].

DISCUSSION

This study highlights the potential of combining computer vision and generative multimodal models to enhance allergen detection and risk evaluation from meal images [6,8,18,22]. YOLOv8n performed well in identifying common allergenic foods, with a solid mean average precision at 50% intersection over union (mAP50), a metric reflecting the model's accuracy in localizing and classifying objects when predicted bounding boxes overlap at least 50% with the ground truth. In this exploratory phase, the model was trained for 10 epochs. Although limited, this training allowed for a preliminary assessment of feasibility. Increasing the number of epochs, together with early stopping to avoid overfitting, could improve performance in future iterations. Gemini 2.0 Flash complemented the pipeline with contextual and multimodal reasoning, adding interpretive depth beyond image recognition. Its transformer-based architecture leverages visual and textual prompts, allowing it to associate visual detections with potential allergenic risks even when they are not visible [4,6,18,22]. However, the system cannot identify invisible or trace allergens, such as those introduced through cross-contamination, and cannot replace full ingredient disclosure or clinical safeguards [9,11,17,25].

One limitation lies in the use of a benchmark dataset instead of real-world images, which limits generalizability. Furthermore, clinical validation in real-life settings is still lacking. Another concern involves ethical and legal risks: if the system misidentifies or fails to flag a relevant allergen, questions of liability arise, especially if the tool is integrated into consumer-facing applications. Developers must also consider informed consent, user education, and clear disclaimers to prevent misuse [11,20,26].

Future directions should include more diverse and representative datasets, clinical testing, and integration with ingredient databases such as OpenFoodFacts to enrich model output. The fast pace of multimodal LLM development — now reaching Gemini 2.5 — suggests that such models may soon outperform task-specific detectors like YOLO, potentially consolidating detection and reasoning into unified AI systems. Despite current limitations, this experimental approach represents a promising step toward practical AI-assisted food safety tools for children and families managing food allergies [8,9,18].

CONCLUSIONS

The presented AI-based prototype, combining YOLOv8n detection and Gemini 2.0 Flash generative inference, achieved high accuracy in allergen identification and produced detailed, contextual risk reports. While not a replacement for medical advice or ingredient transparency, it may serve as a valuable support tool for allergy management in both domestic and public environments. Ongoing development and real-world validation will be essential for safe and effective deployment [8,9,18,22].

Acknowledgements

We gratefully acknowledge the authors of the Allergen30 dataset for providing open access to their data, which was utilized in this study.

Ethical consideration

As this study is based on publicly available online datasets, ethical approval and informed consent are not applicable.

Funding

This research received no external funding .

Conflicts of interest statement

The authors have no competing interests to declare.

Authors' contributions

Methodology, G.M.; writing-original draft preparation, A.P.; data curation, A.S.C.; writing-review and editing, S.C.; supervision, M.M., V.F.; supervision, C.I.; project administration, M.M.d.G. All authors have read and agreed to the published version of the manuscript

Bibliography

1. Gupta RS, Warren CM, Smith BM, et al. The public health impact of parent-reported childhood food allergies in the United States. Pediatrics 2018;142:e20181235. https://doi.org/10.1542/peds.2018-1235
2. Sicherer SH, Sampson HA. Food allergy: Epidemiology, pathogenesis, diagnosis, and treatment. J Allergy Clin Immunol 2014;133:291-307. https://doi.org/10.1016/j.jaci.2013.11.020.
3. Mishra M, Sarkar T, Choudhury T, et al. Allergen30: Detecting food items with possible allergens using deep learning-based computer vision. Food Anal Methods 2022;1-34. https://doi.org/10.1007/s12161-022-02353-9.
4. Ultralytics. YOLOv8 Documentation. https://docs.ultralytics.com/models/yolov8/ (Accessed on: 15/02/2025).
5. Konstantakopoulos FS, Georga EI, Fotiadis DI. A Review of Image-Based Food Recognition and Volume Estimation Artificial Intelligence Systems. IEEE Rev Biomed Eng 2024;17:136-152. https://doi.org/10.1109/RBME.2023.3283149.
6. Google DeepMind – Gemini. https://deepmind.google/technologies/gemini/ (Accessed on: 15/02/2025).
7. Landau T, Gamrasni K, Barlev Y, et al. A machine learning approach for stratifying risk for food allergies utilizing electronic medical record data. Allergy 2024;79:499-502. https://doi.org/10.1111/all.15839.
8. Grabenhenrich LB, Dölle-Bierke S, Worm M, et al. Global trends in anaphylaxis epidemiology and clinical implications. J Allergy Clin Immunol Pract 2020;8:1948-1963.e1. https://doi.org/10.1016/j.jaip.2020.05.041.
9. Allen KJ, Turner PJ, Pawankar R, Taylor S, Sicherer S, Lack G, Rosario N, Ebisawa M, Wong G, Mills ENC, Beyer K, Fiocchi A, Sampson HA. Precautionary labelling of foods for allergen content: are we ready for a global framework? World Allergy Organ J 2014;7:10. https://doi.org/10.1186/1939-4551-7-10.
10. Nwaru BI, Hickstein L, Panesar SS, et al. Prevalence of common food allergies in Europe: A systematic review and meta-analysis. Allergy 2014;69:992-1007. https://doi.org/10.1111/all.12423.
11. Warren CM, Dyer AA, Otto AK, et al. Food Allergy-Related Risk-Taking and Management Behaviors Among Adolescents and Young Adults. J Allergy Clin Immunol Pract. 2017;5:381-390.e13. https://doi.org/10.1016/j.jaip.2016.12.012.
12. Guéant JL, Guéant-Rodriguez RM. Hidden food allergens and the risk of severe reactions in sensitized individuals. J Allergy Clin Immunol 2002;109:1043-1048. https://doi.org/10.1067/mai.2002.122261.
13. Baker MG, Saf S, Tsuang A, Nowak-Wegrzyn A. Hidden allergens in food allergy. Ann Allergy Asthma Immunol 2018;121:285-292. https://doi.org/10.1016/j.anai.2018.05.011.
14. Qian C, Murphy SI, Orsi RH, et al. How Can AI Help Improve Food Safety? Annu Rev Food Sci Technol 2023;14:517-538. https://doi.org/10.1146/annurev-food-060721-013815.
15. Liu D, Zuo E, Wang D, He L, Dong L, Lu X. Deep Learning in Food Image Recognition: a Comprehensive Review. Appl Sci 2025;15:7626. https://doi.org/10.3390/app15147626.
16. Bochkovskiy A, Wang CY, Liao HYM. YOLOv4: Optimal speed and accuracy of object detection. arXiv:2004.10934. https://doi.org/10.48550/arXiv.2004.10934.
17. Skypala IJ, Capucilli P, Wedner HJ. Food-induced anaphylaxis: role of hidden allergens and cofactors. Front Immunol 2019;10:673. https://doi.org/10.3389/fimmu.2019.00673.
18. Yin S, Fu C, Zhao S, Li K, Sun X, Xu T, Chen E. A survey on multimodal large language models. Natl Sci Rev 2024;11:nwae403. https://doi.org/10.1093/nsr/nwae403.
19. Min W, Liu C, Xu L, Jiang S. Applications of knowledge graphs for food science and industry. Patterns 2022;3:100484. https://doi.org/10.1016/j.patter.2022.100484.
20. Ding H, Tian J, Yu W, Wilson DI, Young BR, Cui X, Xin X, Wang Z, Li W. The Application of Artificial Intelligence and Big Data in the Food Industry. Foods 2023;12:4511. https://doi.org/10.3390/foods12244511.
21. Redmon J, Farhadi A. YOLO9000: Better, Faster, Stronger. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 2017, pp. 6517-6525. https://doi.org/10.1109/CVPR.2017.690
22. Brown TB, Mann B, Ryder N, et al. Language Models are Few-Shot Learners. arXiv:2005.14165. https://doi.org/10.48550/arXiv.2005.14165.
23. Gradio. https://www.gradio.app/ (Accessed on: 15/02/2025).
24. Wang L, Niu D, Zhao X, et al. A Comparative Analysis of Novel Deep Learning and Ensemble Learning Models to Predict the Allergenicity of Food Proteins. Foods 2021;10:809. https://doi.org/10.3390/foods10040809
25. Sheth A, Taylor SS, Hourihane JO'B. Deriving individual threshold doses from clinical food challenge data: implications for public policy. J Allergy Clin Immunol 2019;143:2172-2175. https://doi.org/10.1016/j.jaci.2018.12.1024.
26. Li Y, Xu X, Dewey M, et al. Artificial Intelligence in Food Safety: a Decade Review and Current Status. Foods 2023;12:456. https://doi.org/10.3390/foods12030456.