Dean&Francis

ISSN 2959-6157

Automatic recognition of snack packaging time information based on YOLOv11

Kuiyu Chen

Shenzhen University, Shenzhen, Guangdong, China, 518000
Email: 18925266333@163.com

Abstract:

Food safety and information transparency are closely related to national economy and people's livelihood. Against the backdrop of bagged and boxed food dominating the market, accurate and clear labeling and efficient identification of production date and shelf life on packaging are of great significance. However, manual quality inspection has problems such as low efficiency, high cost, and inconsistent standards. The project "Automatic Recognition of Snack Packaging Time Information Based on YOLOv11" aims to solve this problem by focusing on the unique challenges of recognizing snack packaging time information in bags and boxes. Based on advanced YOLOv11 deep learning algorithms, a high-precision and highly robust end-to-end intelligent recognition system is constructed. At the technical level, C3k2 module is used to replace traditional C2f module, and attention mechanism (C2PSA module) is introduced. The detection and recognition module is integrated into a single network head to build an end-to-end pipeline. Hardware technology is integrated to achieve fast and accurate extraction of date information in complex packaging scenarios. In terms of data collection, in response to the problem that the production date and shelf life of most product packaging cannot be in one photo, tools such as Photoshop are used to clip the two in the same image. In addition, the dataset covers various materials and real packaging information with interference such as wrinkles and reflections. Through data augmentation, the original dataset is scaled and rotated to ensure its diversity. The results show that the model can quickly and accurately identify the production date and shelf-life information on food packaging in real-time within 0.5 seconds, and recognize the text of the box selection information through text recognition technology, and finally calculate whether the food has expired. This project implements YOLOv11 architecture improvement and end-to-end lightweight design at the technical level, and has multi scenario adaptability at the application level. Firstly, it can help consumers solve the problem of difficult identification of food packaging information; In addition,

ISSN 2959-6157

this technology is expected to be extended and applied in retail logistics and warehousing scenarios, as well as in the fields of pharmaceuticals and cosmetics.

Keywords: Food packaging, Object detection, YOLOv11, Text recognition, Computer Vision

1. Introduction

1.1 Project Background

Food safety and information transparency are important issues related to national economy and people's livelihood. In the context of bagged and boxed foods (such as snacks, convenience foods, blended drinks, seasonings, dry goods, etc.) dominating the market, accurate and clear labeling and efficient identification of packaging information, especially the production date and expiration date, have become the core link in ensuring food safety, safeguarding consumers' right to know, and ensuring compliant business operations. However, the automated quality inspection of bagged and boxed foods faces unique and severe challenges:

· Low efficiency hinders production capacity: Faced with a high-speed assembly line of hundreds of pieces per minute, it is difficult for human eyes to match the recognition speed.

· High cost and high risk of missed detections: requiring a large amount of manpower for repetitive work, resulting in high costs; Human eyes are prone to fatigue and distraction, leading to a high rate of missed detections. Expired products or packaging with unclear labeling may enter the market.

· Subjectivity leads to inconsistent standards: different quality inspectors have significant differences in the judgment criteria for "clarity" caused by slight wrinkles, ink smudging, local obstruction, etc. on the surface of bags/boxes, resulting in a lack of consistency in quality inspection results.

In this context, there is an urgent need for intelligent quality inspection solutions for bagged and boxed food. This project "Automatic recognition of food packaging production date and shelf life" has emerged. This project focuses on the unique challenges of bag/box packaging and is committed to building a high-precision and highly robust end-to-end intelligent recognition system based on advanced deep learning algorithms (YOLOv11).

1.2 Requirement analysis

1.2.1 Industry demand and policy background

In recent years, consumers have strongly expressed concerns about the production date and shelf life on food labels, which are difficult to find, see, and calculate. These issues directly affect consumers' right to know and choose. As the core carrier of food information transmission, the convenience of identifying production date and shelf life has become an important indicator for measuring consumer experience.

Table 1. Policy Background Related to Food Labeling

Regulatory name	time	specific requirement
Measures for the Supervision and Administration of Food Labeling [1]	2025	Date independent area annotation, black font height ≥ 3.0mm
Announcement on Encouraging Optimization of Label Identification [2]	2024	Encourage enterprises to optimize the location and format of date annotation

On the retail end, supermarkets need to regularly check the shelf life of tens of thousands of products, and manual operations are inefficient and prone to errors. According to the survey, 38% of the delayed removal of expired food is due to manual inspection negligence. In logistics and warehousing scenarios, traditional manual recording methods increase the risk of product expiration by 30%. Food production enterprises are also facing the bottleneck of packaging and printing quality inspection efficiency, and urgently need automation solutions to achieve full chain efficiency improvement.

According to the survey, although existing supermarkets of all sizes have inventory management, they still need to manually check the shelf life or freshness of videos on a regular basis. Some small shops or stores are not equipped with barcode, QR code and other scanning and storage devices, making manual inspection difficult. If an automatic recognition visual system deployed on mobile devices (such as mini programs and apps) is adopted, it can reduce many labor costs.

1.2.2 Scene analysis and user pain points [3]

· Visual complexity: The production date is often printed using inkjet printing, which is susceptible to reflection from packaging materials (such as aluminum cans, vacuum coating), surface deformation (bottle curvature, bag wrinkles), and interference from background patterns. For example, the deformation rate of the spray code on vacuum packaging is high at the wrinkles, which leads to the failure of traditional OCR recognition. At the same time, the unstable lighting environment (low light in warehouses, direct sunlight in supermarkets) further reduces image quality.

· Position uncertainty: Regulations allow dates to be marked on non main display surfaces of packaging, resulting in inconsistent actual position heights. Common labeling locations include "see the seal", "see the bottom of the bottle", "see the side of the packaging bag", etc., which require the system to have full packaging positioning capability.

· Font diversity challenge: Date formats cover various forms such as laser engraving, ink jet printing, and steel stamp pressing, with significant differences in font styles. There are complex situations such as mixed Chinese and English, interference from additional batch numbers, etc., which require the recognition model to have strong noise filtering capabilities.

· Consumer experience dilemma: Most consumers have reported that it is difficult to find the production date, and some packaging may require tools to assist in reading due to small font size or insufficient contrast. Errors in calculating shelf life can also increase the risk of accidentally consuming expired food. According to the survey, most consumers have mistakenly purchased expired products due to calculation errors.

· Pain points in enterprise operation: Obtaining food date information in situations such as expired food promotions and incomplete purchase records of small retailers. Such shops often do not have data managed by large supermarket systems, and the promotion of expired food such as yogurt on supermarket shelves also needs to be managed separately. The use of computer vision recognition technology can better check and update time information.

1.2.3 Existing Technologies and Their Bottlenecks

· Image recognition bottleneck: Existing production date and shelf life recognition technologies mainly rely on OpenCV python and open source OCR technology for image preprocessing and character recognition. However, traditional OpenCV edge detection technology has a high failure rate in recognizing deformed fonts on curved packaging and low light environments.

· Dynamic scene adaptation bottleneck [5]: Existing mobile models (such as YOLOv8n) cannot meet real-time requirements while ensuring recognition accuracy.

Therefore, at this stage, it is necessary to break through the above-mentioned technological bottlenecks and develop a production date and shelf-life automatic recognition model with stronger adaptive features and higher dynamic accuracy.

2. Project Overview

This project proposes an automatic recognition system for the production date and shelf life of food packaging, aiming to solve the problem of date recognition for bagged/boxed foods in high-speed production line scenarios. This project requires the development of an AI vision system based on YOLOv11 deep learning algorithm to achieve high-precision real-time **recognition** and **parsing** of key information such as production date and shelf life on food packaging, replacing traditional manual detection and improving the efficiency of quality control in the food industry. At the technical level, it is necessary to build an end-to-end intelligent recognition pipeline based on the YOLOv11 deep learning detection framework to achieve fast and accurate extraction of date information in complex packaging scenarios. The general technical roadmap is shown in Figure 1.

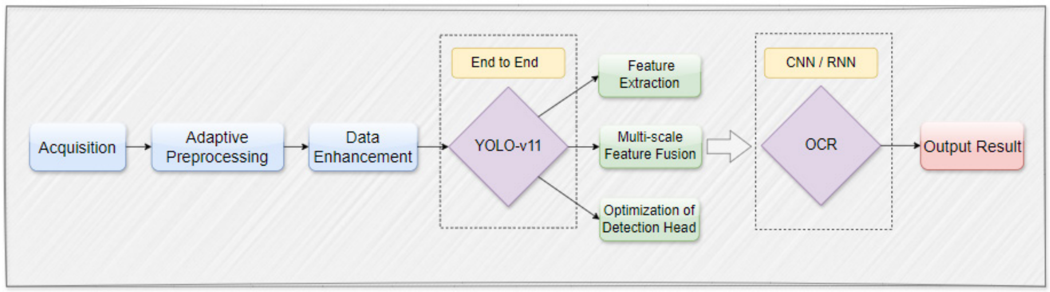

Figure 1. Technical roadmap of the project

The innovation of this project is reflected in two major aspects. In practical applications, this model solves the

ISSN 2959-6157

problem of customers having difficulty distinguishing key information such as production date and shelf life of food packaging. By quickly identifying the location and information of production date and shelf life through cameras and providing feedback to customers, it helps to improve customers' purchasing experience and avoid purchasing expired products. For merchants, it also helps to improve the efficiency of employees in regularly inspecting shelf products. In terms of dataset collection, considering that the production date and shelf life of most product packaging cannot be in one photo, this project uses image processing tools such as Photoshop to clip the production date and shelf life into the same image for data collection; At the same time, the objects of data collection include printed food production information without material interference, food packaging information of different materials with wrinkles and reflections in real photos, and food packaging information of convex bottle bodies in real photos, ensuring that the dataset covers as many packaging types as possible and achieves high robustness.

Expected effect: Place food packaging information on a hardware device equipped with real-time camera detection, adjust the visual system to a clear perspective, and the computer will quickly and automatically select the specific location of the production date and shelf life, recognize textual information, and determine whether it has expired.

3. Technical solution

In response to the core challenges of date recognition for bagged/boxed foods, such as deformation wrinkles, printing defects, background interference, and variable positions, this work builds an end-to-end solution based on YOLOv11. In the official pre training version, YOLO11n version, also known as nano version, is selected. Lightweight models facilitate the implementation of technology on **small devices** and **CPU devices** similar to real-time cameras, as well as fast computation and prediction results. On the model, there are the following features:

· Depthwise Separable Convolutional Kernel: Deep convolution extracts spatial features, while point wise convolution extracts channel features, reducing computational complexity;

· C2PSA [4] attention module: suitable for complex scenes (such as small targets, occlusion);

· Lightweight recognition: integrating recognition modules directly into the detection head to avoid the accumulation of errors in traditional OCR processes;

· Dynamic optimization: Use Adam optimizer [7] to achieve efficient training convergence.

In terms of model training, the datasets used are common real-life food packaging or printed packaging information, and data augmentation is performed through rotation and scaling to improve the fit of the model training. The detailed technical roadmap framework is shown in Figure 2.

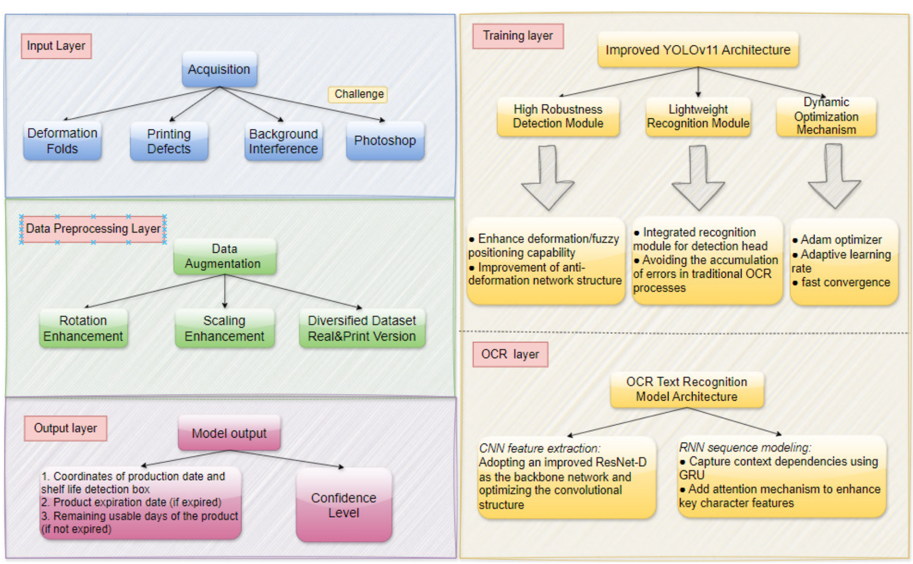

Figure 2. Detailed technical roadmap framework diagram

In addition, YOLOv11 also meets the recognition function of information on food packaging due to its following characteristics.

· Backbone network: YOLOv11 introduces C3k2 module to replace C2f module in previous versions. The C3k2 module is a more efficient computational implementation that crosses stage bottlenecks, using two smaller convolutions instead of one large convolution to accelerate processing speed while maintaining performance;

· Multi scale feature extraction: YOLOv11 retains the Spatial Pyramid Pooling Fast (SPPF) module from the previous version, which facilitates multi-scale feature extraction and enhances the model's detection capability for objects of different sizes.

In the YOLO algorithm, the objectivity score is calculated through a specific formula. The calculation formula for target confidence is as follows:

$$Pr(\text{object}) = \frac{1}{S^2} + \alpha \times IOU_{\text{pred}}^{\text{truth}} \tag{1}$$

Among them, S represents the number of grids that the image is divided into. For each positive sample rectangular box, its Pr (object) value is calculated through annotation information. During the training process, these values will be updated through backpropagation algorithm to enable the model to learn whether each rectangular box contains information about objects.In model training, the loss function is represented by bounding box positioning error and bounding box size error.

$$\lambda_{\text{coord}} \sum_{i=0}^{S^2} \sum_{j=0}^{B} 1_{i,j}^{\text{obj}} \left[\left(x_i - \widehat{x}_i \right)^2 + \left(y_i - \widehat{y}_i \right)^2 \right]^2 \tag{2}$$

$$\lambda_{\text{coord}} \sum_{i=0}^{S^2} \sum_{j=0}^{B} 1_{i,j}^{\text{obj}} \left[\left(\sqrt{w_i} - \sqrt{\widehat{w}_i} \right)^2 + \left(\sqrt{h_i} - \sqrt{\widehat{h}_i} \right)^2 \right]^3 \tag{3}$$

For the character recognition module, Baidu AI Cloud OCR technology is selected as the core solution. OCR technology can convert textual content in images into editable and searchable data, which is crucial for automated processing and analysis of document content. By calling the API interface provided by Baidu AI Cloud, efficient character recognition can be achieved, which not only accelerates the development process, but also ensures the accuracy and reliability of recognition.

Specifically, the extracted time information will be processed to calculate the remaining shelf life of the food. By comparing the current date with the production date and shelf life of the food, it is possible to determine whether the food has expired or is still fresh. This step is particularly important for food safety and quality control, as it can help users make more informed purchasing and consumption decisions.

4. System implementation

4.1 Data Collection and Preprocessing

4.1.1 Data source

Collect packaging samples of common bagged/boxed foods (puffed foods, seasonings, beverages, etc.), covering materials such as plastic composite film, aluminum foil, and laminated paper boxes. As mentioned in Chapter 3, use Photoshop to synthesize images: clip separately captured production dates and shelf life areas into the same packaging image, with a total of 25 packaging samples.

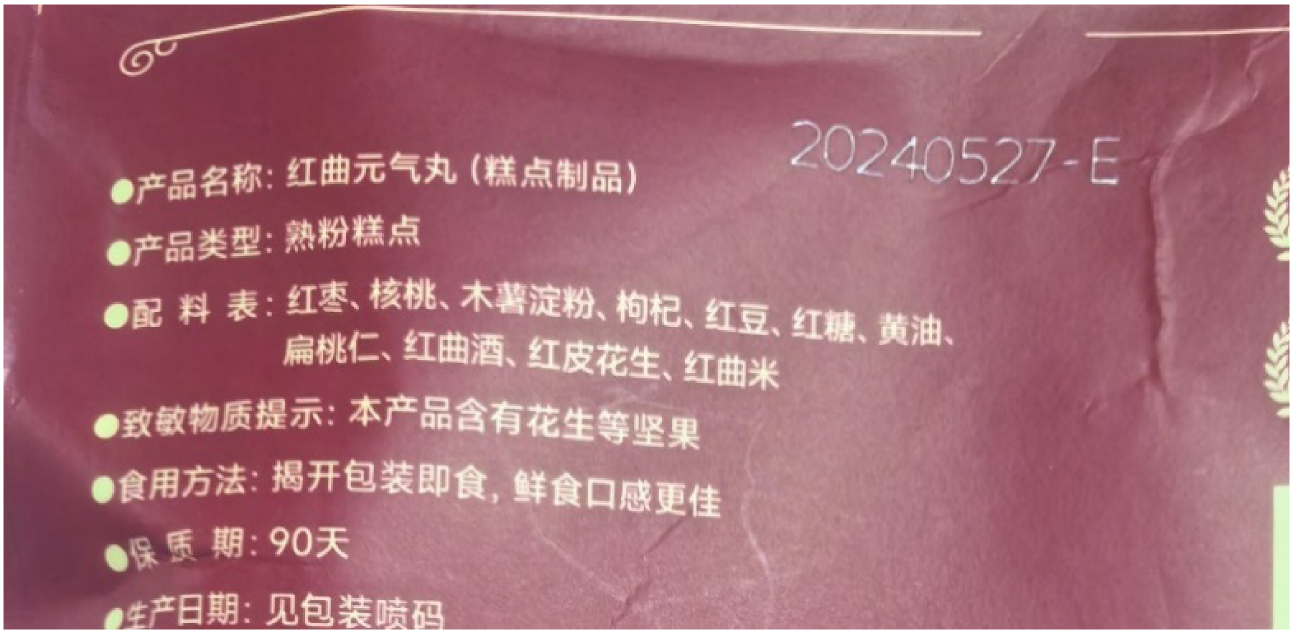

Figure 3. Real shooting picture

ISSN 2959-6157

Figure 4. Photoshop corrected image

4.1.2 Data annotation

Use the 'labelImg' tool to annotate the production date and shelf life areas in the image, and generate corresponding YOLO type annotation files. The annotation file contains the position information (center coordinates), width and height information, and category information (production date or shelf life) of the target object.

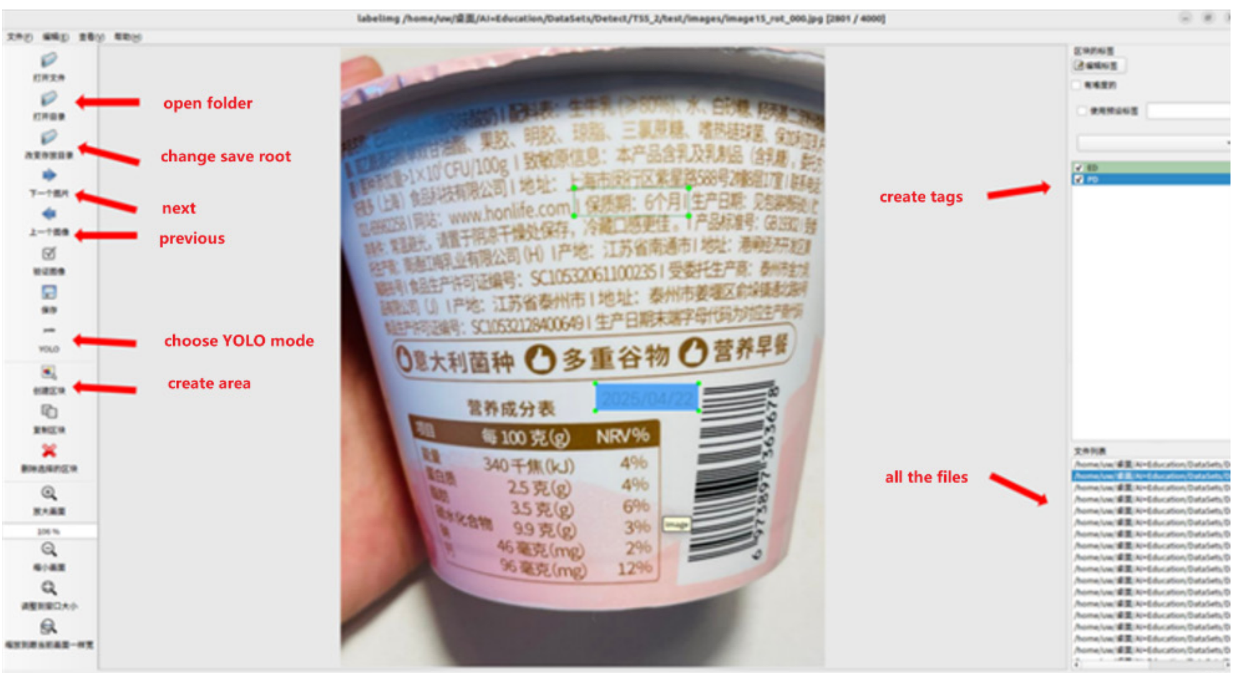

Figure 5. Data annotation using labelImg

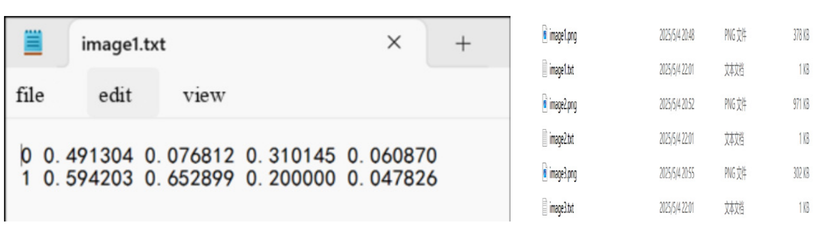

Figure 6. Record the annotation category and position information and save it in txt

4.1.3 Data augmentation [6]

Write a script to apply rotation and random scaling (0.8-1.2

times) enhancement operations to the initial 25 images, increasing the data by 360 times. The script will synchronize the rotation of the images and annotation boxes, ultimately expanding to a **self-made dataset** of 9000 materials.

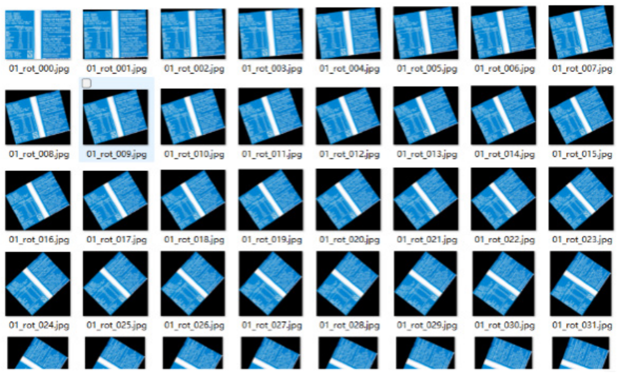

Figure 7. Data Enhancement

4.1.4 Dataset partitioning

Divide the training set, testing set, and validation set according to 8:1:1 to ensure that the model is exposed to sufficient diversity during the training phase, while strictly isolating the test data to objectively evaluate generalization ability. Among them,. yaml specifies the file path, number of yolo tags, and names.

Figure 8. Configuring a self-made dataset

4.2 Model Training and Evaluation

4.2.1 Data loading and preprocessing

Write a data loader to load locally stored datasets into memory and perform preprocessing operations. Preprocessing includes image decoding, resizing, normalization, etc., to make it meet the requirements of the model input. At the same time, the annotation information is parsed and converted to generate the label format required for model training.

4.2.2 Model Training

Input the preprocessed data into the **YOLO11n.pt** model for training. During the training process, the **Adam optimizer** is used to set appropriate learning rates and weight decay parameters to control the training process of the model. Calculate the prediction results of the model through forward propagation, use the loss function to calculate the difference between the predicted results and the true values, and update the model parameters through backpropagation. Real time monitoring of changes in loss values during the training process. When the loss value stabilizes and reaches the preset threshold, the model training is considered complete.

Table 2. Calculation of Partial Losses during Training Process

epoch	time	train/box_loss	train/cls_loss	train/dfl_loss
1	1957.02	1.53214	2.39922	1.38789
2	3898.13	1.16065	1.08879	1.1246
3	5853.6	1.07601	0.8659	1.08498
4	7795.96	1.00589	0.77083	1.05225
5	9745.73	0.91126	0.67022	1.01337

ISSN 2959-6157

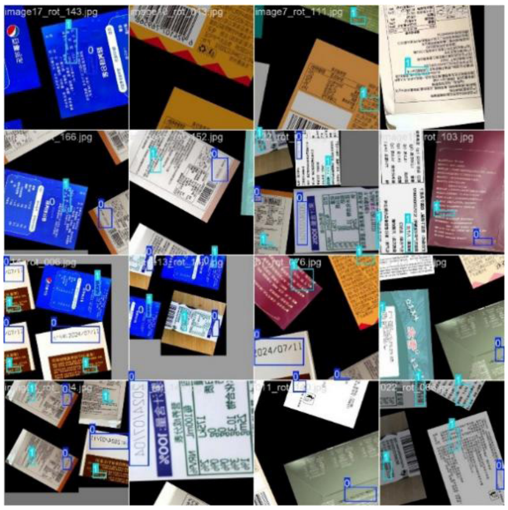

Figure 9. Partial Training Batch

4.2.3 Model Preservation and Reasoning Evaluation

After the model training is completed, save the parameters and structure of the model to a local file. During inference, load the saved model file and restore the model to its trained state for recognition of new images.

The evaluation of model performance mainly consists of two points: first, the accuracy of recognition, which includes the accuracy of box selection information, whether it is close to expectations, and the confidence of labels; The second is the effect of real-time camera recognition, whether the sample can be recognized in a very short time by placing it on the camera.

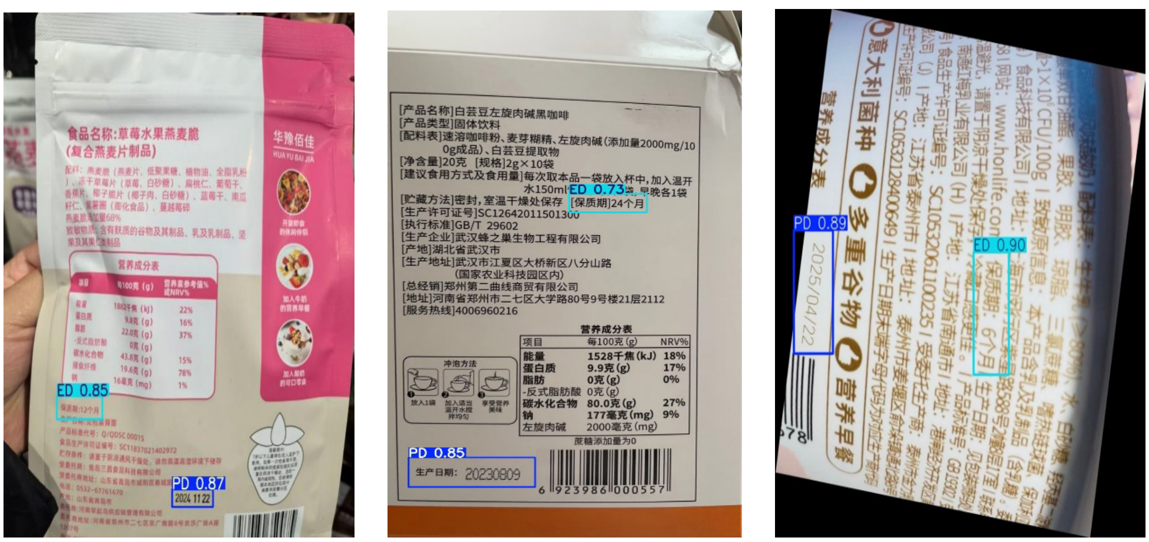

Figure 10. Static Image Prediction

After testing the model, it was found that static photos can accurately frame two key information on food packaging, namely production date and shelf life, with a **confidence level of around 0.8.**

It is worth noting that due to the training of data from **various angles, accurate** and fast recognition can be achieved in real-time prediction regardless of the angle at which the material is placed. During the placement process, if there is a blur, it can even be boxed, thanks to the advantage of small model parameters, which can be **deployed on lightweight devices** for **fast operation**.

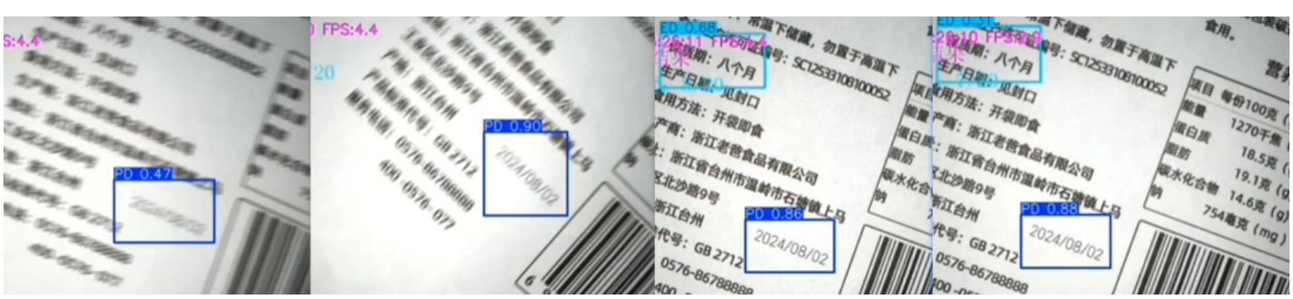

Figure 11. Real time detection

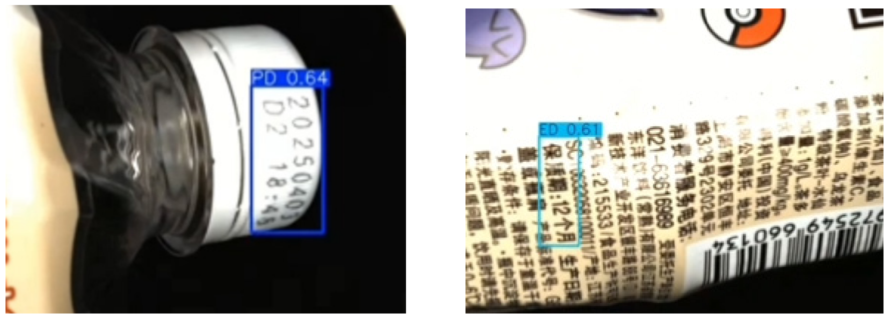

Figure 12. Physical Placement Detection Outside the Training Set

According to the test results, it can be seen that for samples placed at various angles, the corresponding positions of production date and shelf life can be quickly and real-time detected. In addition, under the exposure of special environments such as Figure 12-2, it is also possible to accurately frame the production date and shelf life of untrained beverage bottles, which has high robustness.

4.3 End to end pipeline integration

The end-to-end process refers to the complete automation process from image acquisition input to model inference output recognition results. This project has constructed an efficient and automated food packaging production date and shelf life recognition system by integrating data collection, preprocessing, model inference, and result display. By combining **real-time camera** detection and shooting, this process can directly input the collected images into the model for recognition without manual intervention, thus achieving efficient and high-precision recognition tasks.

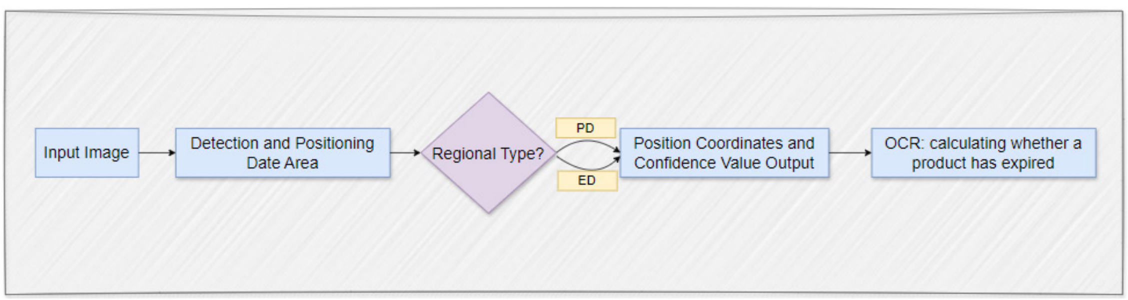

Figure 13. End to end recognition process

4.4 Text recognition and expiration judgment

Extend the text recognition function by calling APIs. Select Baidu Smart Cloud model, apply to AppID, API Key, Secret Key, and use “from aip import AipOcr” in python to call.

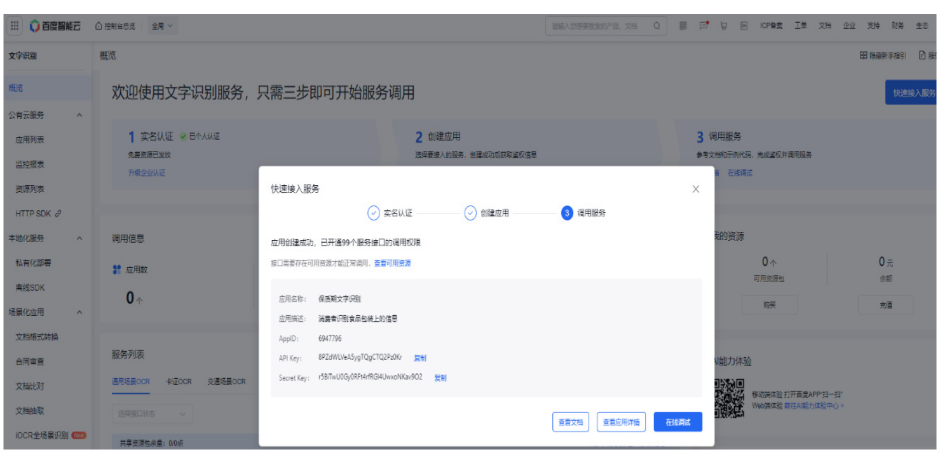

Figure 14. Creating Project Application API

Use the information box obtained from the previous object detection step to select a location, perform text recognition on a small area, and output it in the vicinity. The overall implementation idea of the code is:

ISSN 2959-6157

· Passing API credentials: When creating an OCRService instance, pass in your API key and URL.
· Save temporary file: Save the cropped image area in memory as a temporary file before calling the OCR API.
· Transfer file path: Pass the path of this temporary file to the OCR service.
· Result visualization: Draw the recognized text next to the bounding box of the original image.

After recognizing the text, extract and calculate the text content to determine whether the product has expired. Use regular expressions (RE module) to accurately extract the “production date” and “shelf life” from the OCR recognized string. Extraction of production date:

```
Re.search (r (\ d {4}) [/- year ]\S * (\ d {1,2}) [/- month ]\s*(\d{1,2}), detected_text)
```

The next step is to calculate the date and determine the ‘expiration date’ based on the extracted information. Compare the ‘expiration date’ with the set ‘today’ date (June 1, 2025). The final result is visualized by displaying the judgment results (such as “not expired” or “expired”) and expiration date on the image, and distinguishing them with different colors (**green/red**).

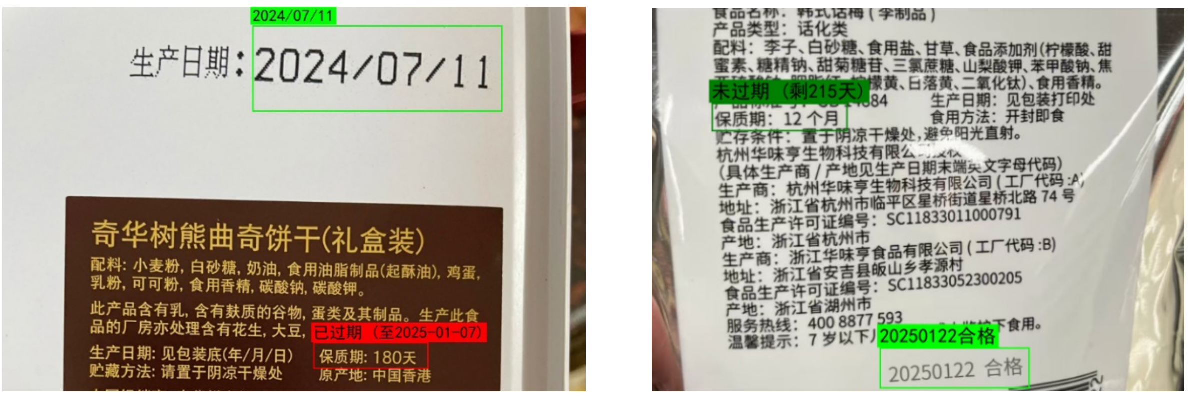

Figure 15. Text recognition and expiration judgment

At the terminal, the system will also provide corresponding outputs, including its **production date, shelf life, expiration date,** and **whether** it has **expired judgment**. For example, in the left figure of Figure 15, the expiration date can be calculated based on its production date and shelf life, and then compared with today’s date to infer whether the product has expired.

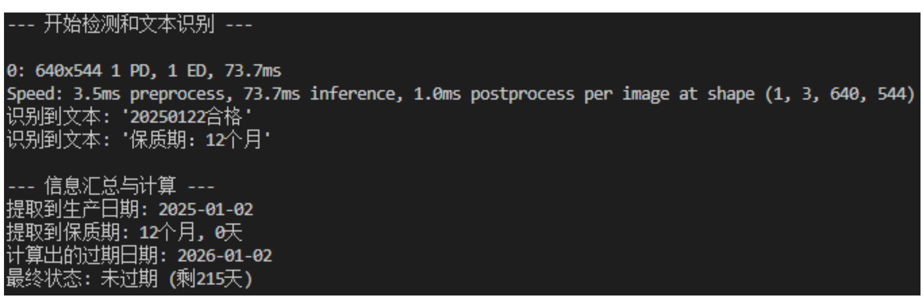

Figure 16. Terminal output result

4.5 Detection+Recognition Fusion

In **4.4**, text recognition for a single detected image has been achieved by adding an OCR module to the annotated boxes detected by YOLO. I now hope to integrate the two tasks and achieve an end-to-end model. The goal is to enable the device to simultaneously detect and identify, and directly provide information on whether it has expired.

Through multi window mode, the left window is used for video stream input, and real-time dynamic monitoring of production date and shelf life annotation boxes is performed, which are selected and displayed in the window. Due to the time required for OCR calling, the right window performs text recognition for the specified frame, followed by determining whether the product has expired. If it has not expired, it will display in green and indicate the remaining time. If it has expired, it will be marked in red and the expiration date will be calculated for display.

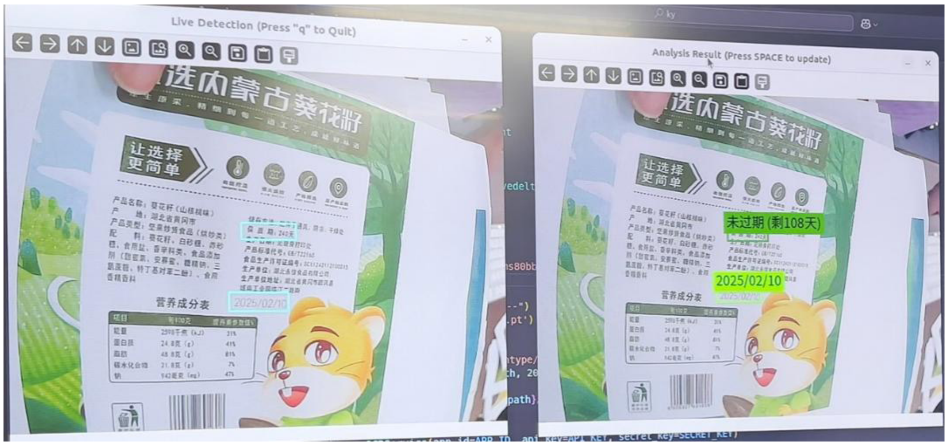

Figure 17. Not expired product

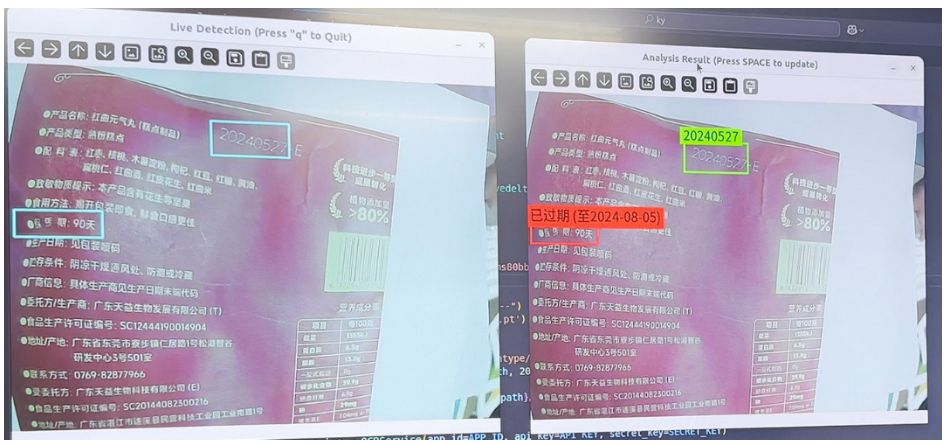

Figure 18. Expired product

4.6 Application Analysis

After incomplete network data analysis and multiple personal tests, the normal time for human eyes to recognize the production date and shelf life of food is about 3-5 seconds, without considering the speed impact caused by fatigue and distraction after long-term and repeated recognition operations. In addition, there are often situations where the shelf life cannot be found after searching for more than ten seconds. This is because the printing font is too small and the printing position is not fixed, and information such as shelf life and production date often appears in the corners or on the other side of the packaging. If this visual model is used, it can achieve fast and accurate real-time recognition within 0.5 seconds, greatly improving detection efficiency and reducing labor costs. If there is no information related to production date and shelf life in the detection area, the model will not provide detection feedback.

This project also tested images outside the dataset and successfully identified the production date and shelf life information on potato chip packaging and beverage bottles outside the dataset images. This indicates that the model has strong generalization ability and has obtained universal rules for capturing shelf life and production date information, avoiding overfitting.

In addition to being used by shop inspectors, it can also be applied from the perspective of consumers to help them quickly find the production date and shelf life of products, thereby determining whether they are fresh. According to multiple tests, the barcode of the product can only scan partial information such as product name, ingredient list, manufacturer, etc., and does not include time information such as production date and shelf life. In addition, some barcodes cannot be scanned except for dedicated systems, making it difficult to quickly obtain information from them. Furthermore, even if the production date and shelf life are scanned, consumers still need to search, extract, and calculate a large amount of information, which can be replaced by this system.

5. Summary and outlook

The main achievements of the project include a **system**

ISSN 2959-6157

that can automatically identify the production date and shelf life of food packaging in real time, and extract information to determine whether it has expired; And a data enhanced **dataset** containing 9000 pieces of food packaging information.

5.1 Summary

This project successfully constructed an automatic recognition system for the production date and shelf life of food packaging, utilizing the YOLOv11 deep learning algorithm to achieve efficient and accurate recognition of the production date and shelf life of food packaging. By using innovative methods such as replacing traditional C2f modules with C3k2 modules and introducing C2PSA attention modules, the model structure is optimized. Combined with a complete data processing flow of data collection, annotation, and enhancement, the model has strong adaptability to complex scenarios. In the experiment, the model achieved high recognition accuracy for static photos with a confidence level of 0.8 to 0.9, and performed well in real-time camera recognition, quickly and accurately identifying food packaging information from different angles with interference. Integrating the detection and recognition module into a single network head avoids the error accumulation problem of traditional OCR assembly lines, and the overall response time of the system is short, meeting the real-time requirements of the production line. By using synthetic data augmentation techniques (such as Photoshop editing production dates and shelf life to the same image), a highly diverse dataset covering 25 types of packaging and 9000 samples is constructed to solve the problem of labeling the separation of two types of information in practical scenarios. Support complex surface recognition for bagged (plastic film, aluminum foil) and boxed (laminated paper box) packaging.

5.2 Outlook

In the future, this project will have vast potential for application expansion. On the one hand, this technology can be integrated into mobile mini programs or small portable devices, making it convenient for consumers to identify food packaging information anytime and anywhere, meeting their concerns about food safety. At the same time, it can also help elderly people in their families easily manage their food and avoid accidentally consuming expired food. On the other hand, the application areas can be expanded to industries such as cosmetics and pharmaceuticals, which also include production date and shelf-life information. By appropriately adjusting and training the model to adapt to the packaging characteristics and information formats of different products, further expanding its technological influence, promoting progress in product quality control and information transparency across multiple industries, providing consumers with more comprehensive and convenient identification services, and enhancing the intelligence level and consumer experience of various industries.

References

[1] State Administration for Market Regulation Measures for the Supervision and Administration of Food Labeling [EB/OL]. (2025-03-14) [2025-06-12]. samr.gov.cn/zw/zfxxgk/fdzdgknr/fgs/art/2025/art_4edcff1e8d894890a012aac1e974c1ff.html.

[2] State Administration for Market Regulation Announcement of the State Administration for Market Regulation on Encouraging Food Enterprises to Optimize the Labeling of Pre packaged Food Production Date and Shelf Life [EB/OL]. (2024-03-26) [2025-06-12]. samr.gov.cn/zw/zfxxgk/fdzdgknr/spscs/art/2024/art_701a6b8b063e4eab85fcb722258c310d.html.

[3] China Science and Technology Network Why is the production date of food and medicine frequently “hidden” due to the blurred printing of location “jumping”. [EB/OL]. (2024-04-15) [2025-06-12]. https://www.stdaily.com/web/gdxw/2024-04/15/content_1957218.html.

[4] Woo, S., Park, J., Lee, JY., Kweon, I.S. (2018). CBAM: Convolutional Block Attention Module. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science(), vol 11211. Springer, Cham.

[5] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934.

[6] Shorten, C., Khoshgoftaar, T.M. A survey on Image Data Augmentation for Deep Learning. J Big Data 6, 60 (2019).

[7] Kingma, D. P., & Ba, J. (2015). Adam: A Method for Stochastic Optimization. International Conference on Learning Representations (ICLR). https://arxiv.org/abs/1412.6980.