Tampere University

Pill Detection

Antti Martikkala 12.12.2024

Euroopan unionin osarahoittama

ETELÄ-POHJANMAAN LIITTO
Regional Council of South Ostrobothnia

Seinäjoki

innokaupungit

Tampere University

Overview of the case study

- Content:
  - Objective: Detect and classify pills using YOLOv8.
  - Motivation: An automated pill detection system can serve as a double-check mechanism, reducing the risk of dispensing errors.
  - Dataset:
    - Size Variability: Pills range from small capsules (e.g., Probiootti) to large tablets (e.g., ParaTabs 1g).
    - Shape Diversity: Includes round, oval, oblong, and capsule forms.
    - Color Range: Mostly white or light-colored pills, with some yellow/orange varieties (e.g., Berex).

Current Challenges:

- Manual verification is prone to human error, especially in high-stress environments like hospitals.
- Pills often look similar in size, shape, and color, making it difficult for healthcare workers to differentiate them visually.

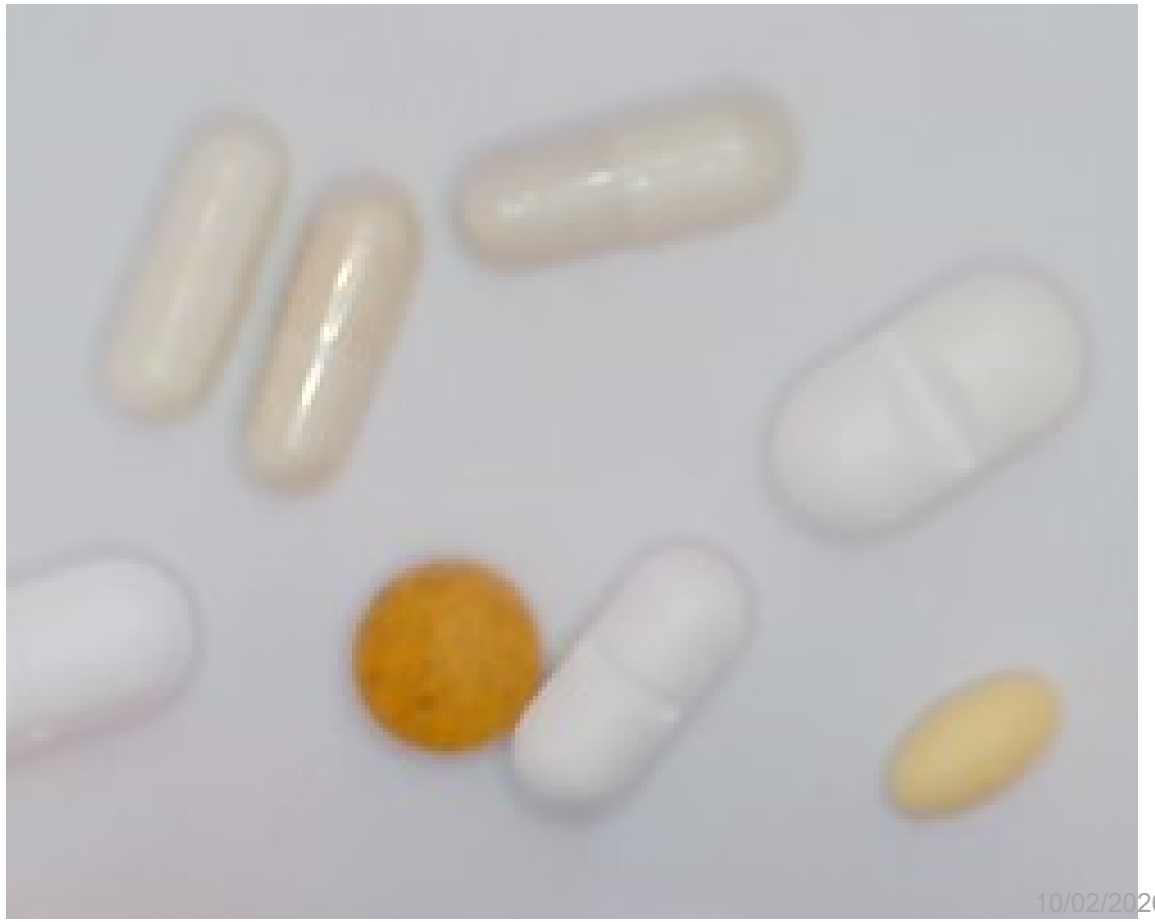

Dataset and Model Training

- Collection of images: Training, validation, and test sets.
- 8 classes of pills: Ibumax400mg, Berex, etc.
- Possible challenges in data: Background clutter, lighting variations, similar-looking pills.
  - Using closed environment with white background and light.
  - A fixed-distance camera setup ensured uniformity in pill sizes across images.
  - Pills were placed at random angles to introduce variability and improve the model's robustness.

Test environment

- USB camera
- Low-cost plant LED
- DNA router box to have constant lighting and fixed camera distance

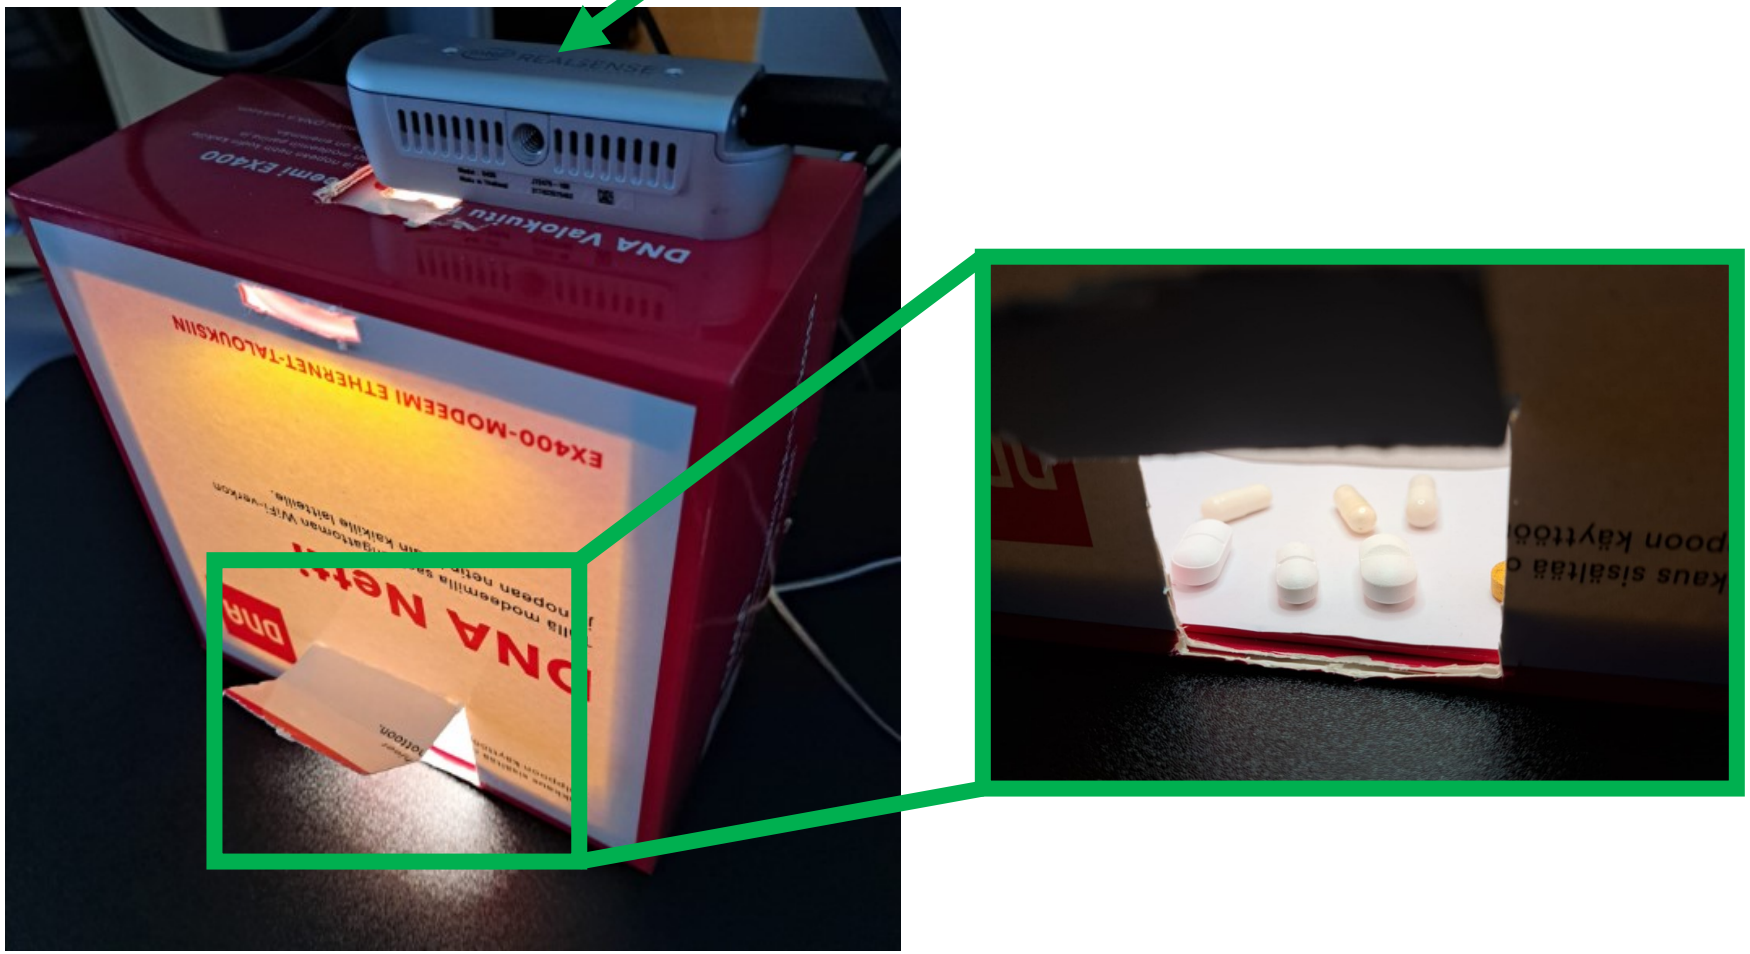

Sample images for the dataset

- Images were annotated with LabelImg
- Classes:
  1. Ibumax400mg
  2. Berex
  3. ParaTabs1g
  4. Somac40mg
  5. Precosa
  6. Probiootti
  7. RSentsyymi
  8. Antepsin

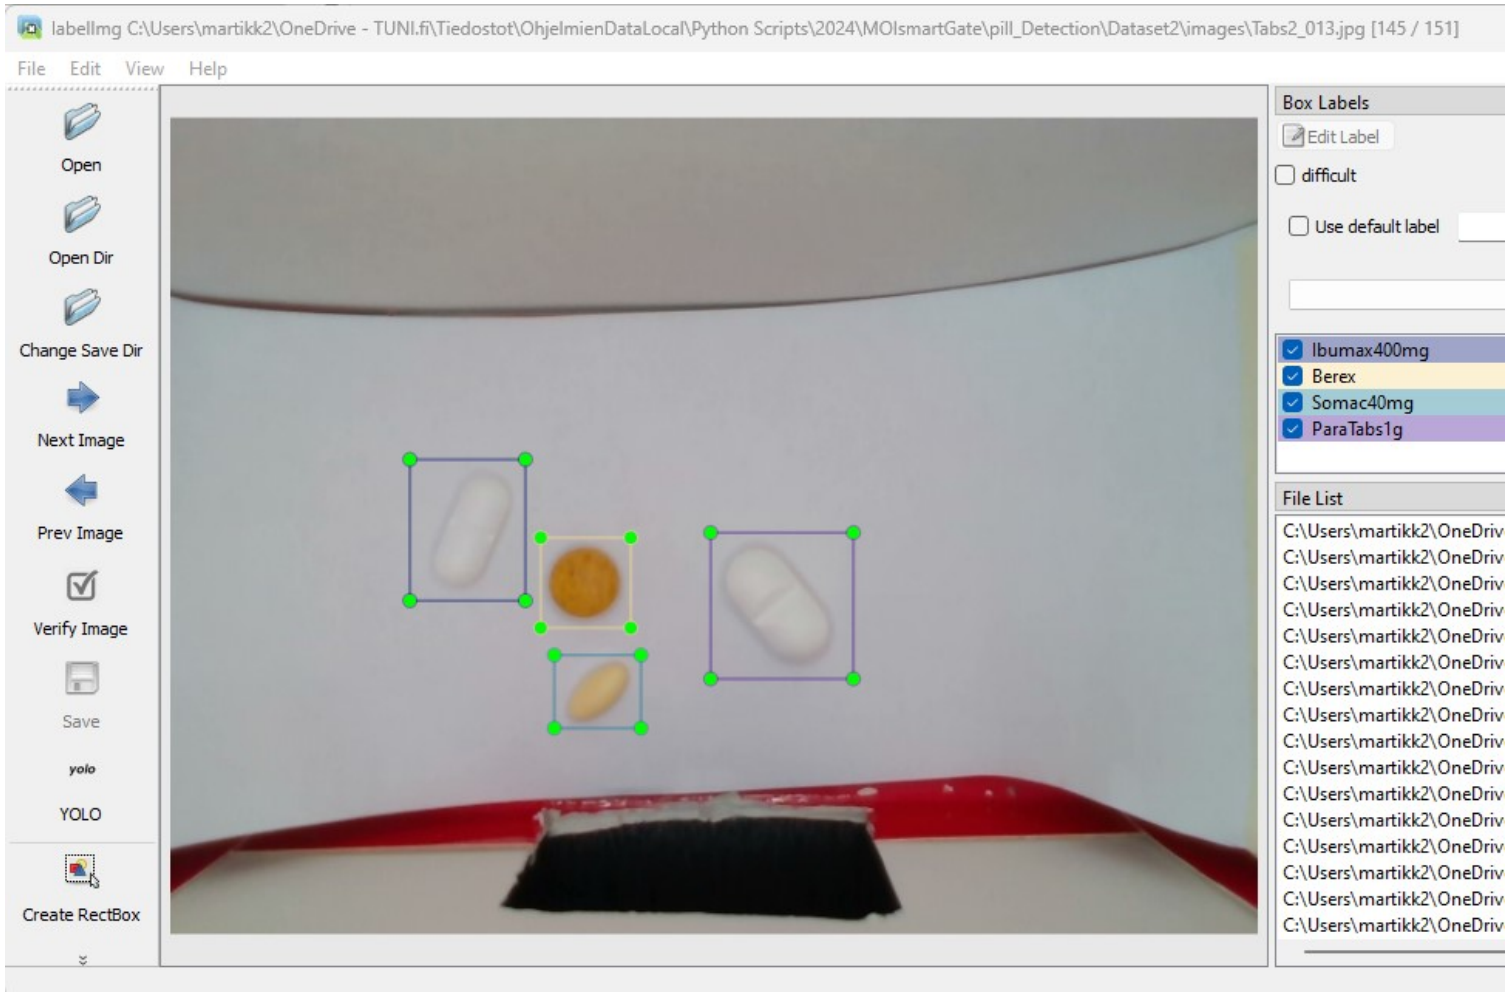

Training Process

- Model:
  - YOLOv8 with pre-trained weights (yolov8n.pt).
  - Fine-tuned for pill detection across 8 classes: Ibumax400mg, Berex, ParaTabs1g, Somac40mg, Precosa, Probiootti, RSentsyymi, and Antepsin.
- Augmentation Techniques:
  - Rotation: Up to ±45° to simulate pills at various angles.
  - Translation: Up to 60% of the image to simulate pills in different positions on the table.
  - Horizontal flipping: 50% probability for symmetric pills.
  - Vertical flipping: 50% probability for vertically symmetric pills.
  - HSV Adjustments:
    - Hue changes: ±10%, Saturation changes: ±20%, Brightness changes: ±20%.
  - Scaling: Disabled to maintain pill size consistency (camera distance is fixed).
  - Shear: Disabled to avoid distortion of pill shapes.
- Training Parameters:
  - Learning rate: Automatically chosen by YOLO's auto optimizer configuration (AdamW with lr=0.000833).
  - Batch size: 16, Image size: 640×640 pixels, Epochs: 50.
- Cache and Optimizer Details:
  - Cache files created for faster dataset loading.
    - Training set: 105 images.
    - Validation set: 30 images.
  - Optimizer: AdamW with lr=0.000833 and momentum of 0.9.
  - Frozen layers: model.22.dfl.conv.weight (deep feature learning layer).

```
Training set class distribution:
Ibumax400mg: 32 images
Berex: 20 images
ParaTabs1g: 37 images
Somac40mg: 12 images
Precosa: 8 images
Probiootti: 8 images
RSentsyymi: 7 images
Antepsin: 31 images

Validation set class distribution:
Ibumax400mg: 7 images
Berex: 8 images
ParaTabs1g: 9 images
Somac40mg: 4 images
Precosa: 3 images
Probiootti: 2 images
RSentsyymi: 2 images
Antepsin: 6 images
```

Confusion Matrix and Metrics

- Precision: 62.4%
- Recall: 82.9%
- mAP50: 71.3%
- mAP50-95: 50.0%
- Underperforming classes (e.g., Probiootti, RSentsyymi).

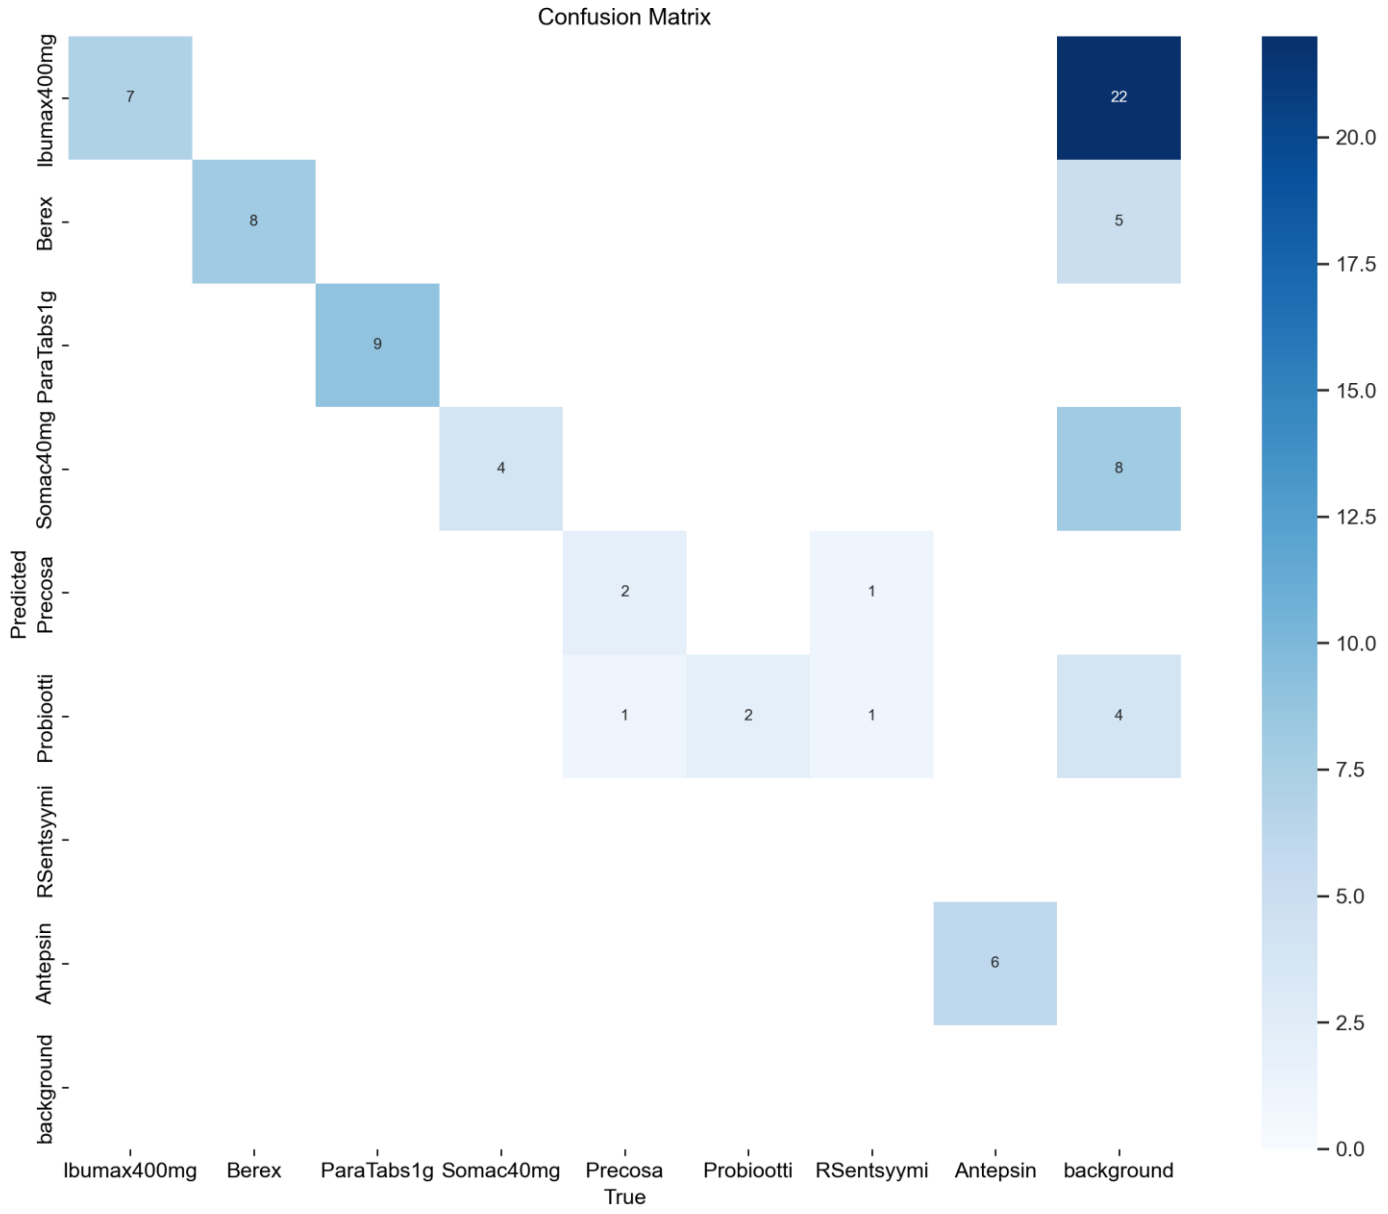

Class-wise mAP:

Some classes (e.g., Probiootti and RSentsyymi) have lower detection performance, while others like Antepsin and ParaTabs1g perform well.

```python
# YOLOv8 Training Log Snippet
Transferred 319/355 items from pretrained weights
Freezing layer 'model.22.dfl.conv.weight'
train: 105 images, 0 corrupt
val: 30 images, 0 corrupt
optimizer: AdamW(lr=0.000833, momentum=0.9)
Starting training for 50 epochs...
Epoch    GPU_mem   box_loss   cls_loss   dfl_loss  Instances   Size
  1/50         0G      3.193      6.617      2.747         14      640
```

These are the per-class mAP values (from the `maps` array):

- `0: Ibumax400mg` : 25.3%
- `1: Berex` : 73.1%
- `2: ParaTabs1g` : 78.7%
- `3: Somac40mg` : 51.1%
- `4: Precosa` : 51.5%
- `5: Probiootti` : 23.9%
- `6: RSentsyymi` : 17.9%
- `7: Antepsin` : **78.5%`

Performance Analysis

- Strengths:
  - High recall (82.9%): Most objects are detected.
  - High mAP50 for certain classes (e.g., ParaTabs1g and Antepsin).
- Weaknesses:
  - Lower precision (62.4%): False positives need to be reduced.
  - Poor performance for specific classes (Probiootti and RSentsyymi), suggesting a lack of sufficient training examples or features for these classes.

Tampere University

Results

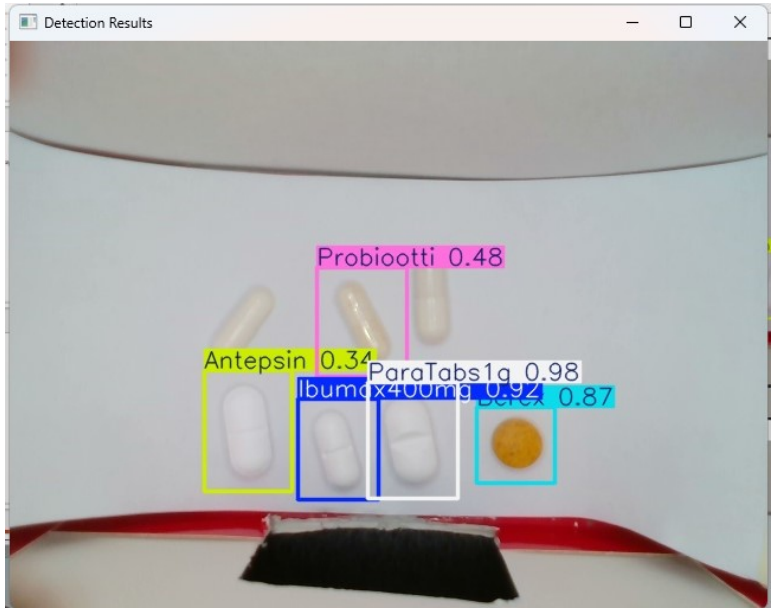

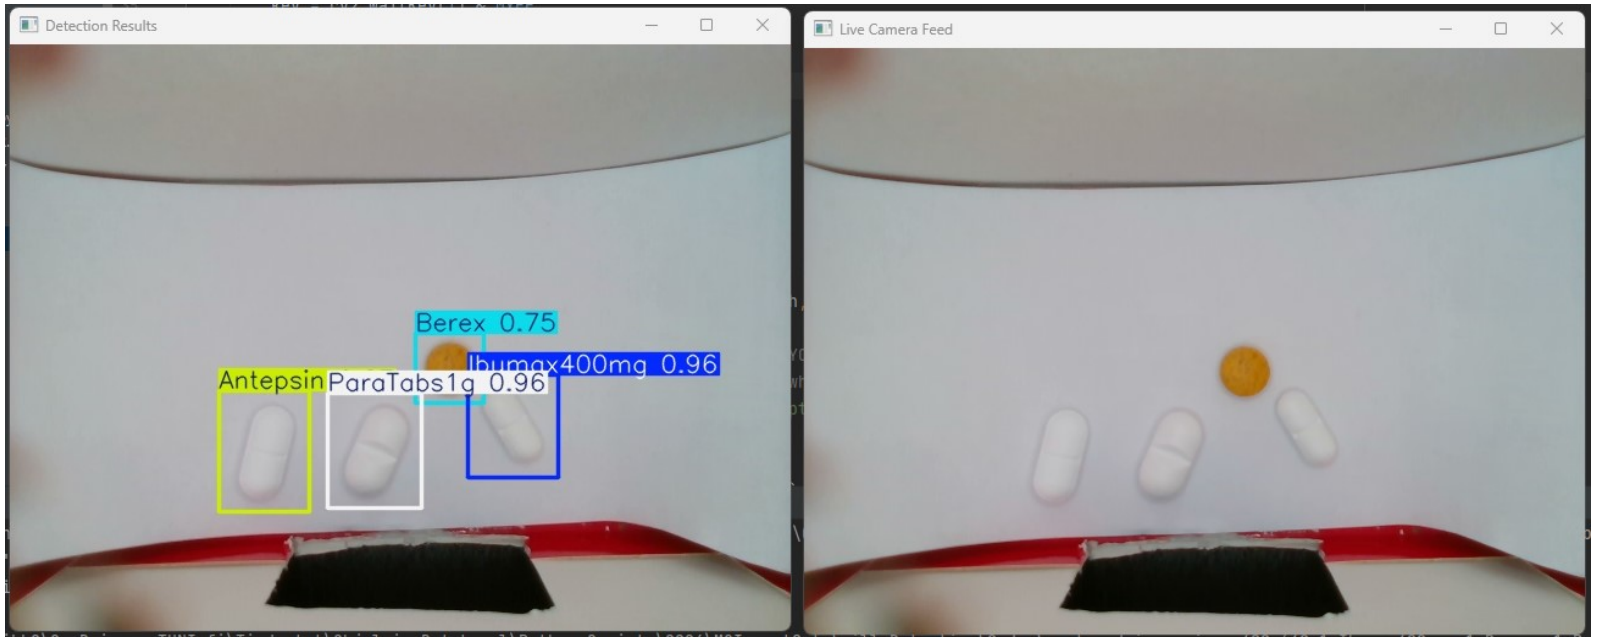

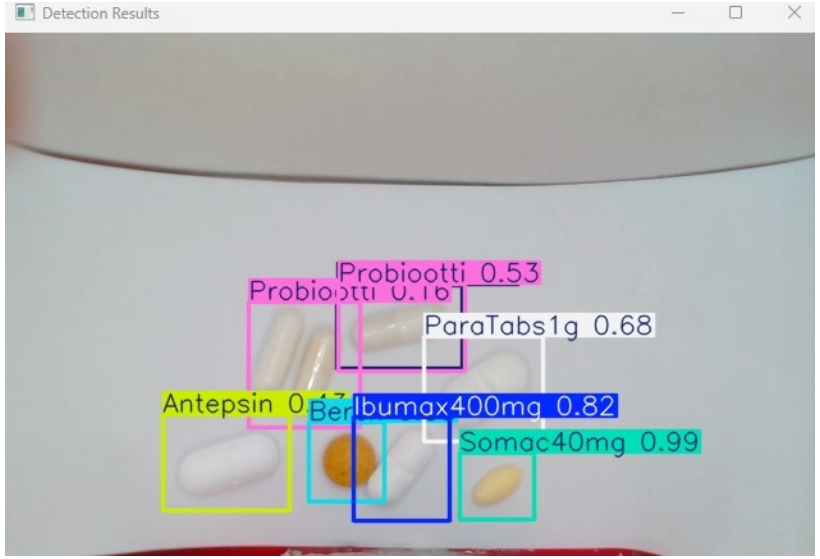

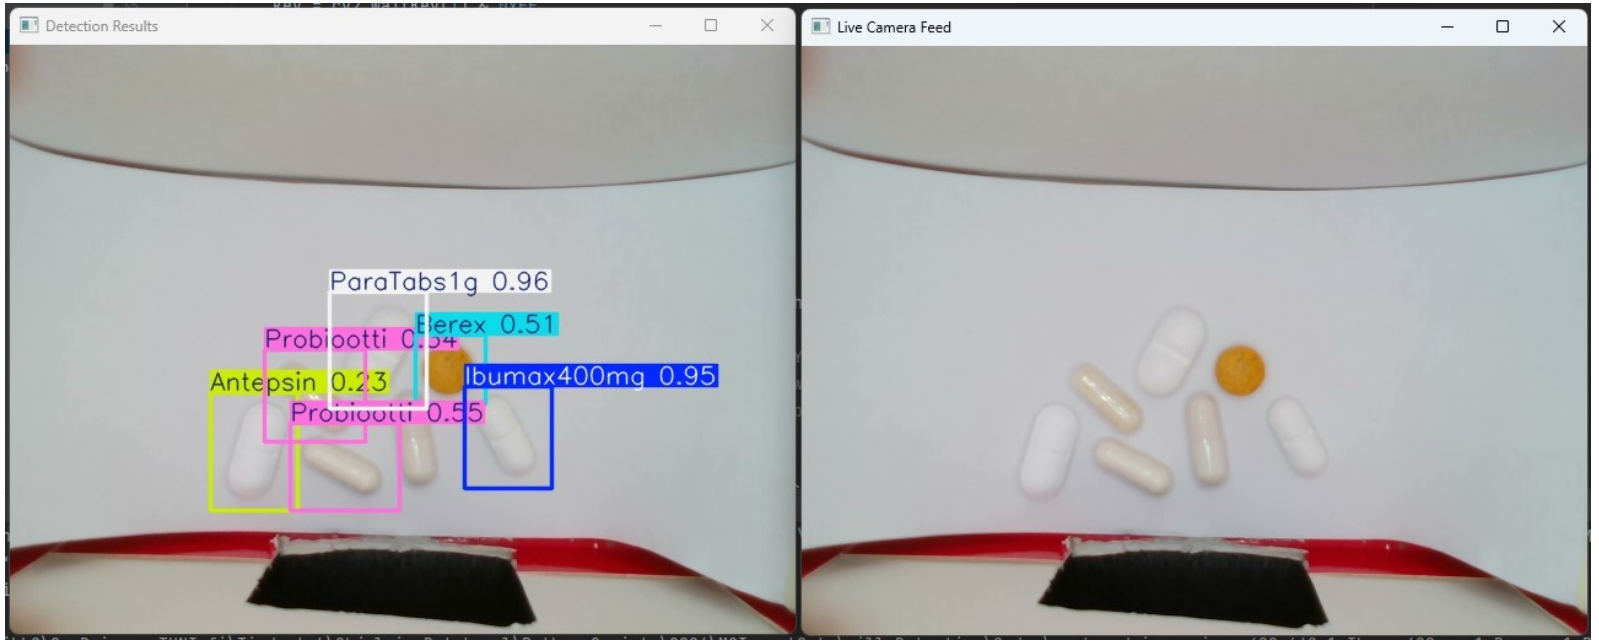

Struggles with similar pills!

Challenges

- Misclassification of similar pills!
- Poor performance for underrepresented classes.
- Sensitivity to lighting and background.
- Limited dataset size.

Future Improvements

- Increase dataset size and diversity.
- More images to dataset of underrepresented classes
- Use better camera
- Use reference information in detection:
  - What pills should be in the picture?
- Explore other model architectures or ensemble methods.
- Real-world deployment and testing.