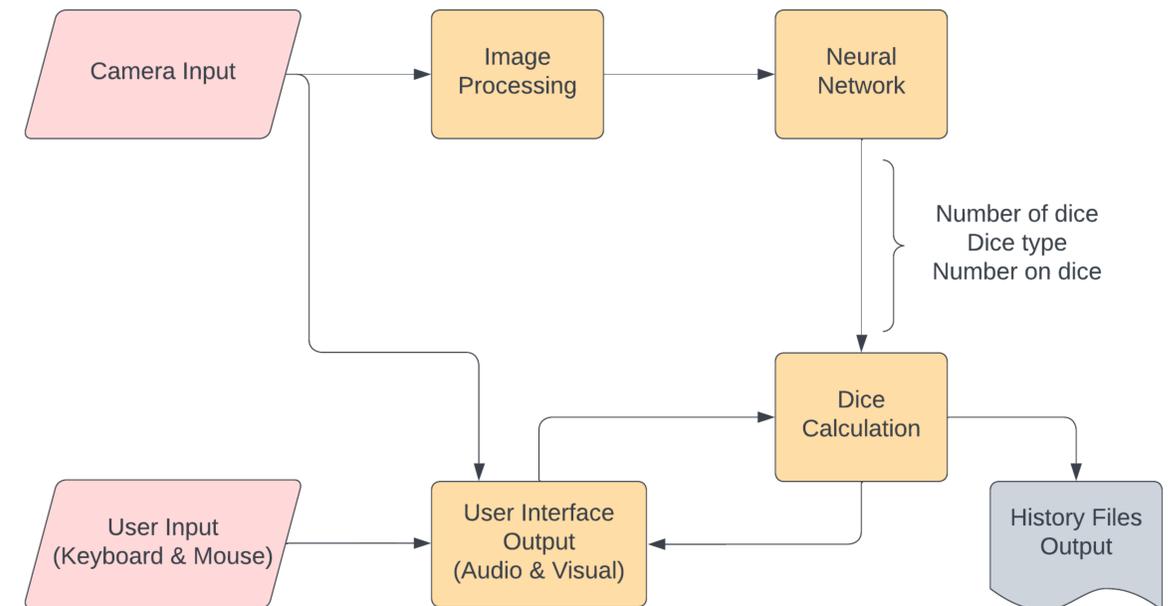


BeholdDie: A Dice Reader

Project

- Tabletop role-playing games (TTRPGs) rely heavily on physical dice, which can pose challenges for many players—especially those with visual impairments or cognitive disabilities. Even for fully able players, fast-paced gameplay can make it difficult to read and calculate dice rolls accurately and efficiently.
- This project developed a real-time dice recognition system using computer vision, designed to identify dice types and rolled values through a webcam. By providing clear visual and auditory feedback, it reduces cognitive load and enhances accessibility during gameplay. Aimed at creating inclusive TTRPG experiences, the system will also support players with learning disabilities by eventually calculating and displaying results automatically, making complex mechanics more approachable for everyone.

System



Methods

We developed a computer vision-based system using YOLOv8 to detect and classify tabletop RPG dice in real time. The process involved:

- Dataset Creation: Collected 14,788 images featuring real-world dice in varied conditions. Around 4,000 images were labeled with bounding boxes using Labellmg, defining 76 unique classes (dice type + value).
- Model Training: Trained a YOLOv8 model using transfer learning, data augmentation, and hyperparameter tuning. Object detection performance was evaluated using mAP, precision, recall, and F1 score.
- UI Design: A Qt-based user interface was built to display detections and apply dice roll calculations. It included visual feedback and groundwork for auditory output, with future accessibility enhancements planned.
- Accessibility Focus: All methods were designed with universal access in mind. Features were developed to reduce barriers for users with visual or cognitive impairments.

Conclusion

The project demonstrated the feasibility of using object detection to recognize and interpret tabletop dice. While the YOLOv8 model achieved a promising mAP@0.5 of 71.7%, performance was limited by the number of labeled examples. Still, precision (85%) and recall (68%) indicate that the system performs reasonably well despite dataset constraints.

The interface provided a functional foundation, and the system shows clear potential to improve accessibility in tabletop gaming. Future work includes full model-UI integration, expanding and labeling the dataset, and refining accessibility features like voice feedback and input-free operation.

This project showed that inclusive, computer vision-based TTRPG tools are not only possible—they're practical and scalable. With continued development, this system could make gaming more welcoming for everyone.