





Product Pitch

We present a guided learning system for Chinese Chess that combines computer vision, a chess engine, and large language models to provide real-time feedback and strategy guidance. The system detects the board state automatically, suggests legal moves using LED indicators, and explains strategies through an AI assistant. Our goal is to transform traditional gameplay into an interactive learning experience for beginners and intermediate players.


System Architecture




System Description

Our system integrates computer vision, a chess engine, and AI coaching into a real-time interactive platform. A YOLOv8 pipeline detects board states, which are validated by a Rust engine and coordinated through a Python state bridge, while a Go pipeline generates AI explanations. A Raspberry Pi-driven LED board provides real-time feedback, and a React interface with a 3D avatar enables user interaction.

System Setup



System Evaluation



Metric	Target	Actual
LED Accuracy	≥ 95%	~98%
LED Latency	< 200ms	~185ms
YOLOv8 Piece Recognition Accuracy	≥ 95%	~99% (vs YOLOX ~87%)
Board Detection Latency	< 200ms	~125ms
LLM Latency	< 2.5s	~500ms (with token limitation)
Overall Latency	< 5s	~2.5s

Conclusions & Additional Information

We successfully achieved our MVP. While we aimed for fully seamless automation, practical constraints required us to introduce design trade-offs. Through this process, we learned the importance of balancing accuracy, latency, and user experience when integrating hardware and software systems. Future work includes further enhancing overall user experience such as adding automated frame capture and piece movement, and improving robustness under diverse environments.