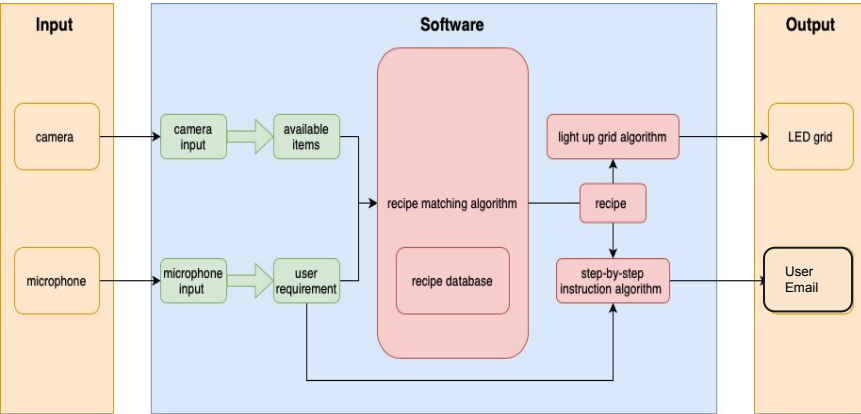


Product Pitch

Our project involves building a smart kitchen assistant that tracks the user's available items in their fridge and based on the available items and user's preference, suggests a specific recipe to the user. SmolKat will then guide the user to different items in the refrigerator that need to be picked out to make the recipe. The full scope of our implementation involves computer vision to track the standard items in the refrigerator (85.7% Accuracy, 1.3s Latency), speech-to-text functionality to interpret user's preference (81% Accuracy, 1.5s Latency), a software recommendation database of recipes and ingredients required (100% Correct Recipe), and a physical guide (LEDs, 100% Correct Grid Lighting) that will direct the user to the appropriate items to complete the recipe.

System Architecture

Implementation Diagram (Hardware & Software)




System Description

Speech Recognition: Convert user speech input to words (requests, like "vegan", "non-dairy")

Image Recognition: Camera takes pictures as input, return recognized ingredient names ("beef", "squash", "banana") and object locations

Recipe Recommendation: Takes the processed requests and ingredients list as input, filter out the appropriate recipes with instructions as output, and send instructions to user email

LED Grids: The grids that holds the corresponding ingredients will light up for user to find them easily



System Evaluation

Requirement	Result	Expectation
Speech Recognition Accuracy	81%	90%
Selected Ingredient Recognition Accuracy	100%(Single)/85.7%(Multiple)	90%
Ingredient Located Grid Correctness	10/10 Tests with Correct Grids Lightened	100% Correctness
Speech Request Latency (95th percentile)	1.489s	2s
Image Recognition Latency (Average)	0.806s(Single)/1.302s(Multiple)	2s
Processing Time for Recommendation	< 0.001s	1s
Correctness of Recommended Recipes	20/20 Correct recommendations	100% Correctness

Trade-off Between Latency and Accuracy For Image Recognition

