

Lab 5 Project Proposal: Sunrise Community Kitchen

Jordan Blake, Alex Chen, Priya Sharma (Instructor Example) | SE 395-2: AI Tools for Social Impact | October 8, 2026

SECTION 1: PARTNER AND PROBLEM

Sunrise Community Kitchen is a three-person nonprofit operating out of a donated church kitchen in the Parramore neighborhood of Orlando. Every Saturday, they distribute groceries and prepared meals to roughly 200 families. Maria Sandoval is the founder, director, and only paid staff member. Dani is a part-time coordinator who manages volunteer communications. A rotating pool of 15 to 20 community volunteers handles the physical distribution each week.

The problem we are addressing is the weekly meal distribution coordination workflow. From Monday through Sunday, Maria manages every step of this process alone: reading and categorizing a multi-page food bank manifest, building a distribution plan that accounts for 200 families with different household sizes and dietary restrictions, writing and sending volunteer instructions, overseeing Saturday distribution, and writing a post-distribution report for the food bank and United Way.

This work takes Maria six to eight hours every week. That does not sound like a lot until you consider what it means when she is unavailable. In February, Maria had the flu. Dani attempted to run distribution without her. She could not interpret Maria's handwritten notes, did not know which families had halal restrictions, and ended up doing first-come-first-served. Three halal families received pork. Two elderly families who always receive home delivery were never called. Maria spent the following week rebuilding relationships.

"If I am out, nothing happens." — Maria Sandoval, Sunrise Community Kitchen

That single sentence is the core of the problem. Sunrise is not understaffed because they cannot recruit volunteers. They are fragile because the knowledge that runs the kitchen exists in one person's head and nowhere else.

SECTION 2: CURRENT STATE

This section documents what actually happens every week, drawn from our Lab 4 workflow analysis.

The Four Steps and Their Time Cost

Step	When	Who	Time	What Happens
Manifest processing		Maria	90 minutes	Maria receives a PDF from the Central Florida Food Bank via email. The PDF lists 200 to 400

	Monday evening, Tuesday			line items with product names, category codes, quantities in cases or pounds, and expiration dates. She prints it, reads every line, and groups items by hand into four categories on a legal pad: protein, produce, dairy, and pantry staples. The PDF format varies week to week. Some weeks it is a clean table. Other weeks it is a raw text dump with inconsistent spacing and abbreviations like "chk thgh" for chicken thighs.
Distribution planning	Tuesday, Wednesday	Maria	2 to 3 hours	Maria cross-references the categorized manifest against her Google Sheet of 200 families. She calculates portion sizes by household (one whole chicken per family of four, two for families of six or more), flags pork products as unsuitable for the 34 halal families, and prioritizes short-shelf-life produce for larger households. She finishes a handwritten distribution list and enters a simplified version into a spreadsheet to print bag labels.
Volunteer coordination	Tuesday through Thursday	Dani, then Maria	45 minutes initial, 1 to 2 hours follow-up	Dani sends a group text to 15 to 20 regular volunteers asking who can work Saturday. Responses come back over 24 to 48 hours as individual texts and occasional emails. Dani tracks responses in a notebook. Maria uses the final confirmed count to finalize station assignments, then writes a run-of-show email from scratch on Thursday: arrival time, station assignments, what to bring, and a summary of that week's items.
Post-distribution reporting	Saturday evening, Sunday	Maria	1 hour	After Saturday distribution ends, Maria hand-counts remaining inventory by visual estimate, photographs leftovers on her phone, and on Sunday writes an email to the food bank and United Way. The report covers: families served, total pounds distributed (manifest total minus estimated leftovers), items not distributed, and any issues such as damaged goods or families turned away. The counts are approximate. Maria acknowledged "close enough" is the standard.

What Works

Two things in the current workflow work well and must not be disrupted.

First, Maria's judgment on priority families. When supply is short, Maria draws on weeks of Saturday conversations and community relationships. She knows the Herrera family's father was deported last month. She knows Mrs. Osei just finished chemotherapy and cannot cook. This knowledge is not in any spreadsheet. It drives the most important decisions in the workflow.

Second, the personal quality of volunteer relationships. Volunteers told us during Lab 4 they keep coming back because "Maria makes you feel like family." The Thursday run-of-show email is not just logistics. It is a personal note to each volunteer. Stripping that away in favor of efficiency risks volunteer retention, which would be a net loss.

Where the Friction Is

- **Single point of failure.** The entire workflow lives in Maria's head. There is no documentation, no backup person who can run distribution independently, and no way to transfer knowledge quickly if Maria is unavailable.
- **Manual manifest processing.** Reading and categorizing 200 to 400 line items by hand every week is the most mechanical part of the workflow and the easiest to automate. Yet it currently takes 90 minutes of Maria's time.
- **Unreliable leftover tally.** Post-distribution inventory counts are visual estimates. If the food bank ever audits the numbers, there could be a discrepancy that damages the relationship.
- **Unknown manifest format.** The food bank's PDF format changes week to week. We do not yet have sample manifests to confirm how variable this is or whether AI parsing would work reliably.

SECTION 3: PROPOSED SOLUTION

We are proposing a Claude-based workflow assistant that handles the mechanical parts of the distribution cycle and presents Maria with organized information so her judgment decisions are faster and better-supported. We are not proposing to replace Maria's decisions. We are proposing to reduce the time she spends on work that does not require her judgment.

What the Solution Does

Manifest categorization. Maria uploads the PDF manifest to Claude. Claude reads each line item and categorizes it using Maria's four-category system (protein, produce, dairy, pantry), not the food bank's storage-based categories. It flags items with expiration dates within five days, notes items that are ambiguous (e.g., plantains appear as "PLT" in some manifests), and outputs a structured table Maria can review on screen or print. Maria reviews the output, corrects misclassifications, and confirms the final list. This step currently takes 90 minutes. We estimate it drops to 15 to 20 minutes of review.

Draft distribution plan with confidence flagging. Claude cross-references the confirmed manifest against the family roster and generates a draft allocation showing what each family is scheduled to receive. When supply is short of demand, it does not make the call. It flags the shortage and shows Maria the relevant context: "38 chickens short of full allocation for this week. The following 12 families received reduced portions last week." Maria reviews the draft and makes all final decisions on priority adjustments. The draft gives her a starting point instead of a blank page.

Volunteer checklist and run-of-show draft. Once Dani's volunteer confirmations are collected, Claude generates a station assignment checklist and drafts the Thursday run-of-show email. The draft pulls the confirmed volunteer names, the week's item summary, and a station map based on available headcount. Maria reviews, personalizes the greeting for each volunteer ("Thanks for coming back, DeShawn" or "Welcome, Lucia"), adjusts any station assignments she wants to change, and sends.

Post-distribution report with manifest-versus-tally cross-check. After Saturday, Maria or Dani enters the remaining inventory count into a simple form. Claude calculates total pounds distributed, families served, utilization rate, and compares it to the previous four weeks. It generates a draft report formatted for the food bank and United Way and drafts the Sunday email. Maria reviews and sends. This step currently takes one hour. We estimate it drops to 15 minutes of review.

Minimum Viable Version

We will not attempt to build all four components at once. The minimum viable version is two components: the manifest parser and the draft distribution plan generator. These are the two most time-consuming steps and the ones with the clearest inputs and outputs. If these two work well and Maria trusts the output, we add the volunteer coordination and report generation components in the second half of the project.

What Stays Human

Four things must remain with Maria and cannot be delegated to the system under any circumstances:

- **Priority decisions when supply is short.** The knowledge driving these decisions lives in relationships and conversations, not in the family roster spreadsheet. AI can present relevant context. Maria makes the call.
- **Cultural and community food knowledge.** Maria knows which families will actually use plantains, taro, or jackfruit. She knows which items will go to waste if allocated to families who cannot use them. No dataset captures this.
- **Final approval on all allocations.** The system generates a draft. Maria approves, adjusts, or overrides it before anything is printed or packed. There is no auto-approve path.
- **Volunteer relationship calls.** The personalized Thursday email and any direct calls to volunteers who are uncertain about Saturday remain Maria's responsibility. The draft is a starting point, not a finished product.

SECTION 4: SDG ALIGNMENT

This project connects most directly to two Sustainable Development Goals.

SDG	Target	Connection	Strength
SDG 2: Zero Hunger	2.1 — Universal access to safe, nutritious, and sufficient food	Sunrise distributes food directly to 200 families each week. The system helps ensure that distribution happens reliably even when Maria is unavailable, which directly supports continuity of food access.	Strong. This is the core mission of the partner.
SDG 10: Reduced Inequalities	10.2 — Empower and promote inclusion of all, irrespective of economic status	The distribution planning component is designed to surface equity considerations: which families received reduced portions last week, which dietary restrictions are not met by available supply. This makes the allocation process more transparent and less dependent on who Maria happens to remember on a given Tuesday.	Moderate. The system can support more equitable allocation but cannot guarantee it. The outcome depends on how Maria uses the information.

We want to be honest about one limitation. The Lab 5 template asks us to identify SDG alignment, and we have done that accurately. But we are not naive about what this project does at scale. Sunrise serves 200 families. The SDG 2 target is global. Our system, if it works perfectly, makes one kitchen more resilient. That matters, but it is a drop in a very large bucket. The SDG framing helps us think about why the work matters. It does not overstate what a semester project can accomplish.

SECTION 5: RISKS AND LIMITATIONS

We identified four categories of risk during our Lab 4 analysis and the Claude stress-test interview. We are being direct about each one because underestimating these risks is how projects fail.

Technical Risk: Manifest Format Variability

This is the highest-risk item in the project. The entire solution depends on Claude being able to read and parse the food bank's PDF manifest. In Lab 4, Maria told us the format varies week to week. Some weeks the PDF is a clean structured table. Other weeks it is a raw text dump with inconsistent spacing, merged cells, and abbreviations. Some weeks she has mentioned the manifest arrives as a scanned image rather than a machine-generated PDF.

If the manifest is a scanned image (not machine-readable text), Claude cannot parse it directly. This is a blocking risk for the manifest categorization component. We would need either an OCR step before parsing or a manual data entry workaround, which may not be practical for Maria each week.

We will not know how serious this risk is until we see actual sample manifests. Getting three to four samples from different weeks is our single most important first action.

Adoption Risk: Rubber-Stamping

The distribution plan component is designed so Maria reviews and approves every AI suggestion before it becomes final. But there is a real risk that after a few weeks of the system working well, Maria starts approving the draft without reading it carefully. This is normal human behavior. It happens with spell-check, with GPS navigation, and with any tool that is usually right.

If Maria rubber-stamps the AI allocation draft without reviewing it, the system creates a false sense of coverage. Errors in the family roster, supply shortages the system miscalculated, or new family crises Maria has not yet logged into any system would all slip through. The output would look reviewed but would not be.

We are building in explicit friction to counter this: the confidence flagging feature is designed to surface cases where the system is uncertain, forcing Maria to read those sections carefully. But we cannot engineer away human attention drift entirely.

Ethical Risk: Encoding Bias in Historical Allocations

If we train the system on Maria's historical allocation decisions, it will learn her patterns. Most of those patterns are good. But some may reflect unconscious biases: families who are better at communicating their needs may consistently receive slightly better allocations than families who are quieter or newer. Families with more complete roster records may be better served than families with sparse data.

We are not building a system that learns from historical allocations in any automated way. The system generates draft allocations from the manifest and roster data, not from past decisions. But we are aware that the family roster itself is a biased document. Families with complete records (household size, dietary restrictions, notes fully filled in) will receive more accurate AI-generated allocations than families with sparse records, who tend to be the newest and often most vulnerable families.

We will document this limitation in the handoff package and recommend that Maria prioritize updating roster records for new families.

Scope Limitation: What We Are Not Building

Three things are out of scope for this project, and we are stating them explicitly to manage expectations:

- **A full volunteer communication platform.** Volunteers still contact Dani by text and email. We are not building a scheduling app, a volunteer portal, or a two-way communication system. We are helping organize the responses Dani already collects.
- **Dietary preference matching at the individual item level.** The system will flag items unsuitable for halal families and flag nut allergy concerns. It will not build a personalized nutrition

profile for each family or optimize allocations for health outcomes. That is a much larger system requiring data we do not have.

- **Integration with the food bank's systems.** We have no access to the food bank's inventory system. We work from the PDF manifest they send. If the food bank changes their PDF format significantly, the manifest parser may need to be updated.

SECTION 6: WHAT WE NEED FROM THE PARTNER

The following list is prioritized. The first two items are required before we can build anything. The third and fourth items are needed for the second half of the project.

Required (Weeks 8 to 9)

1. **Three to four sample manifest PDFs from different weeks.** We need at least one clean week and one messy week. If any recent manifests arrived as scanned images rather than machine-generated PDFs, we need a sample of those too. Without these samples, we cannot test whether Claude can parse the manifest at all. This is the gating item for the entire project.
2. **Thirty minutes watching Maria build a distribution plan on a real week's data.** Not a demonstration, not a walkthrough after the fact. We need to sit with Maria on a Tuesday or Wednesday while she is actually working through the manifest and making allocation decisions. We want to see the moment she flags a shortage, the moment she overrides a portion size, and the moment she makes a judgment call that is not in the spreadsheet. We cannot design the confidence flagging feature without knowing what Maria pays attention to when she second-guesses herself.

Secondary (Weeks 10 to 11)

- **An anonymized version of the family roster spreadsheet.** We do not need real names or addresses. We need the structure: household size ranges, the distribution of dietary restrictions, and examples of the notes column so we can test how the system handles unstructured free text. We will confirm with Maria and the instructor how to handle this appropriately given privacy considerations.
- **A sample Saturday tally sheet.** The post-distribution inventory count Maria currently does by hand. Even a photo of one recent week's legal pad tally would help us design the report generation component. We need to understand what information she captures and in what format before we ask her to enter it differently.

SECTION 7: TIMELINE AND MILESTONES

Weeks	Milestone	Deliverable
8 to 9	Manifest parsing test	Get sample manifests from Maria. Test Claude parsing on each sample. Document where parsing works cleanly, where it struggles, and whether

scanned PDFs are a blocking issue. Present findings to partner and agree on whether to proceed with the manifest component or design a workaround.

8 to 9	Draft allocation prototype	Build a basic version of the distribution plan generator using one week's anonymized roster data and a parsed manifest. Test whether the draft allocation is useful enough for Maria to review rather than rebuild from scratch. Get her honest feedback on whether the output format matches how she thinks.
10 to 11	Real-week test	Run the manifest parser and draft allocation generator on a real week's data alongside Maria's manual process. Compare outputs. Document every case where Maria's judgment differed from the AI draft and ask her to explain why. Iterate on the confidence flagging based on what we learn.
12 to 13	Report generation and full workflow test	Add the post-distribution report component. Run all four components on a full week from manifest through Sunday report. Document total time saved and any points where Maria needed to intervene. Begin drafting the backup operator guide.
14 to 15	Handoff preparation	Train Maria on the full workflow. Train one backup person (Dani or a senior volunteer) using the documentation. Run a test where the backup person uses the system without Maria present and assess whether distribution could have functioned. Finalize all handoff materials.
16	Final showcase	Present to panel. Demonstrate the system live or via walkthrough video. Share time savings data, Maria's feedback, and honest assessment of what worked, what did not, and what the kitchen would need to sustain this after the semester ends.

We are building in buffer at weeks 10 to 11 because the real-week test is the most likely place to surface problems we did not anticipate. If the manifest parsing is unreliable or Maria's feedback on the draft allocation requires us to redesign the output format, we need time to absorb that without losing the second half of the project.

SECTION 8: SUCCESS CRITERIA

We are proposing four success criteria. Each one is specific and observable. None of them are "the system worked" or "the partner was satisfied," which are too vague to be useful.

Criterion	What We Will Measure	How
Time reduction	Maria's weekly coordination time	Maria logs her time on distribution coordination tasks for two weeks before the system is introduced and two weeks after. We compare the before and after totals. We are targeting a

	drops from 6 to 8 hours to 3 to 4 hours	reduction to 3 to 4 hours, not to zero. Tasks that require her judgment still take real time.
Override transparency	Maria can identify which AI suggestions she overrode each week and explain why	The distribution plan interface shows Maria which items she changed from the AI draft before finalizing. At the end of each test week, we ask her to walk us through her overrides. If she cannot remember why she changed something, the override tracking is not working. If she can describe her reasoning for each change, the system is keeping her in the decision loop.
Backup operator readiness	A backup person can run the workflow using the documentation without Maria present	In weeks 14 to 15, we ask Dani or a senior volunteer to use the system and the handoff guide to build a distribution plan for a real week's data without asking Maria for help. We observe where they get stuck, where the documentation is unclear, and whether the output they produce is close enough to what Maria would have produced. This is the real test of whether we solved the single-point-of-failure problem.
Partner endorsement	Maria says "I would use this next week"	We ask Maria this question directly at the end of week 15, after the backup test but before the showcase. Not "Is this useful?" or "Is this impressive?" but "Would you actually use this next week without us here?" A yes means the system is practical. A no with explanation is the most valuable feedback we can receive before the semester ends. We will report both the answer and what it took to get there.

One criterion we are intentionally not including: the number of families served. Sunrise serves roughly 200 families every week regardless of how much time Maria spends on coordination. Improving the workflow does not change that number. Claiming SDG 2 impact on families served would be misleading. The real impact is on organizational resilience and Maria's capacity.

"The goal is not to save Maria time. The goal is to make the kitchen survive a week when Maria is sick." — Team Jordan Blake, Alex Chen, Priya Sharma

If Dani can run a Saturday distribution using the system's output as a starting point, with reasonable judgment calls and no major failures, the project succeeded. That outcome is achievable by December. It is worth building.