Deep Research, SDG Mapping, and Bias Detection: AI Conversation Logs

Lab 2 AI Conversation Logs

Student	John Doe (Instructor Example)
Date	September 22, 2026
Tools	Gemini Deep Research, Perplexity, Claude

This document contains my full conversation logs from the three AI tools used in this lab, organized by step. Each section shows my prompts and the tool's responses, lightly condensed where the output was very long. I have not edited the AI responses except to shorten them for length. Yellow boxes are my notes to myself written during the lab.

STEP 2: GEMINI DEEP RESEARCH (PRIMARY RESEARCH ENGINE)

Sub-Query 1 of 3: Food Bank AI Adoption

MY PROMPT

```
I am a university student researching how AI is being used in food banks and food
distribution nonprofits in the United States. I need a comprehensive research report
covering:

1. Real examples of organizations in this space using AI today (with names, not
hypotheticals)
2. What specific AI applications have evidence of working vs. what is still
experimental
3. Peer-reviewed studies, reports, or case studies documenting outcomes
4. What the main risks and limitations are for AI in this domain
5. Which applications are accessible to organizations with fewer than 200 staff
members

For every claim, include the source. Distinguish between peer-reviewed research,
organizational reports, news coverage, and marketing material. If something is
unverified, say so.
```

GEMINI DEEP RESEARCH RESPONSE (SUMMARIZED)

Research Plan: I'll search for information about AI adoption in US food banks across four areas: inventory management, demand forecasting, donor engagement, and logistics optimization.

Demand Forecasting and Inventory Management

Feeding America has partnered with Google.org on an AI-powered tool called MealConnect that uses machine learning to match food donations with local agencies in real time. The system processes data on food availability, agency capacity, and transportation logistics to optimize distribution (Feeding America, 2023).

The Greater Chicago Food Depository implemented a demand forecasting model developed with Northwestern University's Kellogg School of Management that reduced food waste by approximately 15% in its first year of operation (Orgut et al., 2022).

A 2023 study published in the European Journal of Operational Research examined AI-driven inventory management at food banks and found that machine learning models outperformed traditional statistical methods for perishable food allocation, reducing spoilage by 12-18% across three pilot sites (Balcik et al., 2023).

Logistics Optimization

Second Harvest Heartland (Minnesota) partnered with Microsoft's AI for Good program to develop route optimization algorithms, reportedly cutting fuel costs by 20% (Microsoft AI for Good, 2022).

Google DeepMind has worked with the World Food Programme to develop AI models for supply chain optimization in humanitarian contexts. The underlying logistics algorithms have been adapted for domestic food banks through the Feeding America network (DeepMind Blog, 2023).

Donor Engagement

The Food Bank of the Rockies implemented an AI-powered donor segmentation system using Salesforce Einstein that increased recurring donations by 22% over 18 months (Salesforce.org Case Study, 2023).

Copia, a VC-backed food recovery platform, uses predictive analytics to match food donors with recipients and claims to have diverted over 10 million pounds of food from landfills (Copia website, 2024).

Challenges and Gaps

- Most food banks lack the technical staff to implement or maintain AI systems independently
- Data quality remains a significant barrier
- A 2024 report by the Urban Institute noted that only 8% of US food banks had adopted any form of AI or advanced analytics
- Privacy concerns around client data limit the types of AI applications food banks can deploy

Lots of impressive-sounding claims. The DeepMind-WFP claim and the 8% statistic both feel like they need checking. Also noticed that the Salesforce and Microsoft examples come from those companies' own marketing. Flagging all of these for Perplexity verification.

Sub-Query 2 of 3: SDG Connections

MY PROMPT

```
I need to map the work of Second Harvest Food Bank of Central Florida, a nonprofit
that distributes over 100 million meals annually through 800+ partner agencies
across six Florida counties, to the UN Sustainable Development Goals. Do not just
list which SDGs are "relevant." I need:

1. Specific SDG targets (e.g., Target 11.4, not just "SDG 11") that connect to their
actual operations
2. Evidence that this type of work contributes to measurable progress on those
targets
3. Examples of similar organizations that have documented SDG-aligned outcomes
4. Which SDG connections are strong (backed by evidence) and which are a stretch

Cite your sources. If a connection is based on a published framework or report, name
it.
```

GEMINI DEEP RESEARCH RESPONSE (SUMMARIZED)

> **SDG 2 (Zero Hunger)**
>
> - **Target 2.1** (end hunger and ensure access to safe, nutritious food year-round): AI demand forecasting directly supports this target by reducing gaps between food availability and need. Predictive models allow food banks to anticipate spikes in demand during specific periods (Gundersen et al., 2021).
> - **Target 2.4** (sustainable food production systems): AI-driven logistics reduce food waste in the supply chain, keeping edible food out of landfills.
>
> **SDG 12 (Responsible Consumption and Production)**
>
> - **Target 12.3** (halve per capita global food waste by 2030): AI-powered matching systems like MealConnect reduce food waste by connecting surplus food with agencies that can distribute it before spoilage. ReFED estimates improved food rescue could divert 2.5 million tons annually (ReFED, 2024).
>
> **SDG 1 (No Poverty)**
>
> - **Target 1.2** (reduce poverty in all dimensions): Food assistance frees household income for other needs. A USDA study found that SNAP recipients who also received food bank assistance spent an average of $180 more per month on non-food essentials (USDA ERS, 2022).
>
> **SDG 3 (Good Health and Well-Being)**
>
> - **Target 3.4** (reduce premature mortality from NCDs): Nutritious food access reduces diet-related disease risk. Several studies link food insecurity to higher rates of diabetes and cardiovascular disease (Seligman et al., 2010).

> **SDG 4 (Quality Education)**
>
> - **Target 4.2** (ensure access to quality early childhood development): Research shows food insecurity impairs children's cognitive development and school performance (Jyoti et al., 2005).
>
> **SDG 10 (Reduced Inequalities)**
>
> - **Target 10.2** (promote social and economic inclusion): Food banks serve marginalized communities disproportionately affected by food insecurity.
>
> **SDG 13 (Climate Action)**
>
> - **Target 13.2** (integrate climate change measures): Food rescue reduces methane emissions from landfilled food waste.
>
> **SDG 17 (Partnerships for the Goals)**
>
> - **Target 17.17** (encourage effective partnerships): The food bank AI ecosystem depends on cross-sector partnerships between tech companies, nonprofits, universities, and government agencies.
>
> *[Response continued with academic mapping literature. Key source: Sachs et al. (2022), Nature Sustainability. Truncated for length.]*

> *Gemini mapped Second Harvest to 8 SDG targets. That feels like too many. Some of these are real stretches. SDG 4 and SDG 13 connections are plausible but very indirect. SDG 10 has no specific evidence. I need to cut this down to the ones I can actually back with evidence. The lab instructions say "fewer strong connections are better than many weak ones."*

Sub-Query 3 of 3: Tool Landscape

MY PROMPT

```
I am looking for AI tools that a mid-size food bank (50-200 staff, limited tech
budget) could realistically adopt. I need:

1. Specific tool names with pricing (free tier, monthly cost, enterprise pricing)
2. What data each tool collects and where it is stored
3. Who built the tool (startup, Big Tech, open source community)
4. Whether an open-source alternative exists for each paid tool
5. Any documented case of a food bank or similar nonprofit successfully adopting
each tool

Focus on tools that exist today and have real users, not products that are "coming
soon" or in beta with no public access.
```

GEMINI DEEP RESEARCH RESPONSE (SUMMARIZED)

CRM and Donor Management

- **Salesforce Nonprofit Cloud:** $60/user/month (nonprofit pricing). CRM with Einstein AI add-on for donor segmentation and predictive giving. All data stored on Salesforce servers. Built by Salesforce ($31B revenue). Open alternative: CiviCRM (free, self-hosted).
- **Bloomerang:** Starting at $119/month. Donor retention analytics with some AI features. Built by Bloomerang (Indianapolis-based startup). No direct open-source alternative.

Food Rescue and Logistics

- **Copia:** Custom pricing (enterprise). Predictive analytics for food waste reduction and surplus matching. VC-backed startup. Data includes surplus volumes, partner info, logistics. No direct open-source equivalent.
- **MealConnect (Feeding America):** Free for Feeding America member agencies. Matching platform for food donations. Google.org partnership. Available only within Feeding America network.
- **Spoiler Alert:** Custom pricing. AI-powered food waste management platform. Acquired by Lineage Logistics in 2022.

General Nonprofit AI Tools

- **Google Analytics:** Free. Web traffic analysis for donation pages. Data stored on Google servers.
- **Canva (with AI features):** Free tier available, $12.99/month for teams. AI-powered design for marketing materials.
- **ChatGPT / Claude:** Free tiers available. General-purpose AI for grant writing, communications, data analysis. $20/month for paid tiers.

[Response continued with implementation considerations and case studies. Truncated for length.]

Good overview but I notice the most detailed examples (Salesforce, Copia) are the most expensive and most likely to create vendor dependency. The free tools (Google Analytics, ChatGPT) get less attention even though they are probably more realistic for Second Harvest. Need to run the corporate capture filter on Salesforce and Copia in Step 4.

STEP 3: PERPLEXITY (SOURCE VERIFICATION)

Verification 1: DeepMind-WFP Partnership

MY PROMPT

```
Did Google DeepMind partner with the World Food Programme (WFP) to develop AI models
for supply chain optimization? I found this claim in an AI research report but want
to verify it with primary sources. Please cite your sources with URLs.
```

PERPLEXITY RESPONSE (SUMMARIZED)

> I could not find evidence of a direct partnership between Google DeepMind and the World Food Programme for supply chain optimization.
>
> - **Google.org and WFP:** Google.org (Google's philanthropic arm, not DeepMind) has provided grants to the WFP Innovation Accelerator. This is a funding relationship, not a joint AI development project.
> - **Palantir and WFP:** The WFP partnered with Palantir Technologies in 2019 on a supply chain data project, which drew criticism from privacy advocates.
> - **DeepMind's published work:** DeepMind has published research on logistics optimization, but I found no publications or press releases about food bank or WFP partnerships.
>
> It is possible that the claim conflated Google.org's grants to WFP with DeepMind's logistics research. These are separate entities within Alphabet.

> *FABRICATED. Gemini conflated Google.org's WFP grants with DeepMind's logistics research. Flagging in my source evaluation table. This is the kind of hallucination the lab warned us about. Not including this in my final report.*

Verification 2: MealConnect AI Claims

MY PROMPT

```
Does Feeding America operate an AI-powered tool called MealConnect? What does it
actually do? Is there evidence that it uses machine learning for matching food
donations with agencies? Please cite sources.
```

PERPLEXITY RESPONSE (SUMMARIZED)

> Yes, MealConnect is a real platform operated by Feeding America, but the AI claims need qualification.
>
> MealConnect is a technology platform that connects food donors with local food banks and agencies. It was launched in 2015 and expanded nationally with support from Google.org.
>
> Feeding America describes it as a "technology platform" that "simplifies the process of donating food" rather than as an AI tool. A 2021 Google.org blog post describes the partnership as providing "technical volunteers and funding" but does not specify machine learning as a component.
>
> **Bottom line:** MealConnect exists and works as a matching platform, but describing it as "AI-powered" may overstate its current capabilities. Algorithm does not equal AI.

> *EXAGGERATED. MealConnect is real, but calling it "AI-powered" is an overstatement. Feeding America's own language says "technology platform," not "AI." I will include this in my report but note the distinction.*

Verification 3: Urban Institute 8% Statistic

MY PROMPT

```
Did the Urban Institute publish a 2024 report stating that only 8% of US food banks had adopted AI or advanced analytics? Please verify with a direct source.
```

PERPLEXITY RESPONSE (SUMMARIZED)

> I could not find a 2024 Urban Institute report making this specific claim about food bank AI adoption rates.
>
> There is a related finding: The NTEN & Heller Consulting 2024 Nonprofit Digital Investments Report found that over 60% of nonprofits now use AI in some capacity, a very different number from 8%. The figure is also skewed toward large organizations. The 8% figure may be extrapolated or fabricated.

> *CANNOT VERIFY. The 8% statistic may be fabricated. The NTEN 2024 report found 60%+ for all nonprofits, a completely different number. I will not cite the 8% figure. Using the NTEN & Heller Consulting (2024) finding instead.*

Verification 4: Dutch Childcare Benefits Scandal

MY PROMPT

```
Verify: Did the Dutch government resign in January 2021 over an algorithmic childcare benefits scandal that affected over 26,000 families? Was it called the "toeslagenaffaire"?
```

PERPLEXITY RESPONSE (SUMMARIZED)

> Yes, this is verified and well-documented.
>
> The Dutch childcare benefits scandal (toeslagenaffaire) led to the resignation of Prime Minister Mark Rutte's government on January 15, 2021. A parliamentary inquiry found that the Dutch tax authority used an algorithm that disproportionately targeted families with dual nationality, flagging them as fraudulent.
>
> Approximately 26,000 families were wrongly accused and forced to repay benefits. Amnesty International characterized the scandal as a violation of the right to social security and the right to non-discrimination.

> *FULLY VERIFIED. All details check out. Strong example for my bias audit section.*

Verification 5: Salesforce Einstein at Food Bank of the Rockies

MY PROMPT

```
Did Food Bank of the Rockies implement Salesforce Einstein for donor segmentation,
resulting in a 22% increase in recurring donations? Please verify with primary
sources.
```

PERPLEXITY RESPONSE (SUMMARIZED)

> I found partial evidence for this claim.
>
> Salesforce.org does publish case studies about nonprofit customers, and Food Bank of the Rockies is listed as a Salesforce customer. However, I could not find a specific case study citing the "22% increase in recurring donations" figure or confirming Einstein AI specifically.
>
> Food Bank of the Rockies does use Salesforce for CRM (confirmed on their careers page). Whether they use the Einstein AI add-on specifically is not publicly confirmed.
>
> The 22% figure may come from a Salesforce marketing document rather than an independent evaluation.

> *SELF-REPORTED. They use Salesforce, but the 22% figure is likely vendor marketing, not independent evaluation. Good example of corporate marketing disguised as evidence. Rating credibility as "Low" in my table.*

Verification 6: Copia's 10 Million Pounds Claim

MY PROMPT

```
Copia claims to have "diverted over 10 million pounds of food" using predictive
analytics. Can you verify this figure from an independent source, not Copia's own
website?
```

PERPLEXITY RESPONSE (SUMMARIZED)

> Copia exists and operates as a food recovery technology platform based in the Bay Area. The company uses predictive analytics to match food surplus with nonprofit recipients.
>
> The "10 million pounds" figure appears on Copia's own website and in their press materials. I could not find an independent audit or third-party verification of this number. Copia has received coverage in Fast Company, TechCrunch, and Forbes, but these articles cite Copia's own reported figures without independent verification.

> The company has received VC funding from investors including Andreessen Horowitz and Y Combinator, but investor backing does not verify operational claims.

> *SELF-REPORTED. The platform exists and operates, but the 10M figure has no independent verification. Media coverage just repeats their claims. I will note this as "self-reported" in my source evaluation table.*

Verification 7: 1.3 Billion Tons Food Waste Statistic

MY PROMPT

```
I keep seeing the claim that "1.3 billion tons of food is wasted globally each year"
attributed to the UN FAO. Is this figure current? What is the most up-to-date data?
```

PERPLEXITY RESPONSE (SUMMARIZED)

> The 1.3 billion tons figure comes from the FAO's 2011 report "Global Food Losses and Food Waste." It is over a decade old.
>
> The most current data comes from UNEP's Food Waste Index Report 2024, which uses different methodology and estimates 1.05 billion tonnes of food waste at the household level alone (not including retail and food service waste). The methodologies are not directly comparable.
>
> Many AI research reports and nonprofit websites still cite the 2011 figure because it is more widely known, even though it is outdated.

> *OUTDATED. The 1.3B figure is from 2011. Current UNEP data uses different methodology. I will use the UNEP 2024 figure and note that the older number is still commonly cited but outdated.*

STEP 4: CLAUDE (BIAS DETECTION AND SYNTHESIS)

Prompt 1: Bias Detection

MY PROMPT

```
I am researching how AI tools can help food banks like Second Harvest Food Bank of
Central Florida. My Gemini Deep Research gave me optimistic findings about AI
adoption in food banks, and my Perplexity verification caught several problems: a
fabricated DeepMind-WFP partnership, an overstated "AI-powered" label for
MealConnect, an unverifiable Urban Institute statistic, a vendor case study passed
off as evidence, and an outdated food waste figure still being cited as current.

Based on these patterns, analyze the following biases in my research so far:
```

```
1. Geographic bias: Are these findings mostly from the US/Europe? What perspectives
from the Global South are missing?
2. Data representation: Whose data trained the models that generated this research?
Who is overrepresented?
3. Corporate framing: Which of these findings come from companies selling AI tools?
How does that framing differ from what independent researchers would say?
4. SDG framing bias: Does mapping everything to the SDGs privilege certain kinds of
work?
5. Missing voices: If I interviewed the actual staff at Second Harvest, what would
they say that none of these AI sources capture?

Be direct. I need honest critique, not encouragement.
```

CLAUDE RESPONSE (SUMMARIZED)

> **1. Geographic and Demographic Bias**
>
> Every example in your research comes from US-based organizations and English-language sources. Gemini returned nothing from Latin America, Sub-Saharan Africa, or South Asia, where food insecurity is most severe. The AI solutions being discussed are designed for well-resourced American food banks with digital infrastructure, internet connectivity, and staff who can operate software platforms. These do not transfer to contexts where the problem is most urgent. Any recommendation you make is drawn from a sample biased toward organizations that already have resources.
>
> **2. Survivor Bias**
>
> Every AI implementation example was a success story. No tool surfaced a food bank that tried AI and abandoned it, or one where an algorithm produced harmful recommendations. Even though the NTEN 2024 report says over 60% of nonprofits use AI in some capacity, that adoption is concentrated in large, well-funded organizations. Small rural food pantries, church-based food closets, community fridges, and Indigenous food sovereignty programs serve some of the most vulnerable populations and have zero AI capacity. Organizations that fail with AI rarely document the failure publicly.
>
> **3. Corporate Framing Bias**
>
> - **MealConnect:** Feeding America's own language calls it a "technology platform." Gemini upgraded it to "AI-powered." This is definitional inflation that makes AI seem more pervasive than it is.
> - **Salesforce Einstein:** A vendor case study claiming 22% improvement with no independent verification.
> - **Microsoft AI for Good:** The "20% fuel cost reduction" comes from Microsoft's own reporting.
>
> **4. SDG Framing Bias**
>
> The SDG framework privileges measurable, quantifiable outcomes. It makes food distribution visible (tons diverted, meals served) but obscures the relational and community-building work that food banks do. A volunteer who drives meals to a homebound senior and checks on them weekly is doing work that matters but does not map neatly to any SDG target.

> **5. Missing Voices**
>
> If you interviewed Second Harvest staff, they would likely talk about: the emotional toll of the work, the gap between what funders want to hear and daily reality, the challenge of retaining volunteers, the frustration of navigating government food safety regulations, and the relationships with partner agencies that no technology can replace. None of this shows up in AI-generated research because it is tacit knowledge held by practitioners, not published data.

The survivor bias point is the insight I would not have reached on my own. I was focused on the organizations that use AI and never asked about the much larger number that do not. The SDG framing bias about relational work being invisible is also something I want to include.

Prompt 2: Corporate Capture Filter

MY PROMPT

```
Look at the AI tools I identified in my research: Salesforce Nonprofit Cloud, Copia,
MealConnect, Bloomerang, Spoiler Alert, Google Analytics, and ChatGPT/Claude.

For each tool, answer:
1. Who built it and what is their business model?
2. What data does the tool collect, and does the company use that data for anything
beyond the stated purpose?
3. What happens to the organization's data and workflows if this vendor raises
prices or shuts down?
4. Is this tool making the organization dependent on a platform they cannot leave?
5. Does an open-source or self-hosted alternative exist?

This is the "corporate capture" test from Iazzolino and Stremlau applied to real
tools.
```

CLAUDE RESPONSE (SUMMARIZED)

> **Salesforce Nonprofit Cloud**
>
> - **Builder:** Salesforce ($31B revenue). Business model: subscription SaaS with annual contracts.
> - **Data risk:** All donor and client data stored on Salesforce servers. Terms of service allow aggregated data use for product improvement. Donor giving patterns become training data for Salesforce's AI models.
> - **Lock-in:** High. Salesforce migration is expensive and complex. Nonprofit pricing increases over time. Organizations build custom workflows that are difficult to port.
> - **Open alternative:** CiviCRM (open source, self-hosted). Less polished but the organization owns its data completely.

> **Copia**
>
> - **Builder:** VC-backed startup (Andreessen Horowitz, Y Combinator). Business model: enterprise SaaS.
> - **Data risk:** Collects surplus food data, logistics data, partner information. Data ownership terms unclear in public-facing materials.
> - **Lock-in:** Medium. If Copia shuts down (common for VC-backed startups), the food bank loses its logistics platform entirely.
> - **Open alternative:** No direct equivalent. Google Sheets plus Google Maps API for basic matching is a low-tech workaround.
>
> **MealConnect**
>
> - **Builder:** Feeding America with Google.org support. Available only to Feeding America member agencies.
> - **Data risk:** Lower than commercial tools since Feeding America is a nonprofit. But Google.org involvement means some data flows through Google infrastructure.
> - **Lock-in:** Low for data, but high for network access. Leaving Feeding America means losing MealConnect access.
>
> **Google Analytics, ChatGPT/Claude**
>
> - **Risk:** Lower because they are general-purpose tools. Switching costs are minimal. But conversation data with ChatGPT and Claude may be used for model training unless opted out.
>
> **Corporate capture pattern:** When Google, Microsoft, and Salesforce donate AI tools to food banks, they gain: (1) access to distribution data, (2) platform dependency, (3) marketing case studies, and (4) a seat at the policy table. As Iazzolino and Stremlau argue, this positions tech companies as essential partners while extracting data and creating relationships the nonprofit cannot easily exit.

> *This is exactly what the lab is asking for. The Salesforce lock-in analysis is something I would not have thought to look at. The Copia VC shutdown risk is real. Including the Salesforce and Copia analyses in my corporate capture table in the report.*

STEP 5: CLAUDE (CROSS-REFERENCING ASSISTANCE)

Prompt 3: Building the Source Evaluation Table

MY PROMPT

```
I need to build a source evaluation table for my report. Here are all the sources
that came up across my three Gemini reports and my Perplexity verification:

- Feeding America Annual Report (2025)
- UN FAO State of Food Security (2024)
- Sengul Orgut & Lodree (2023), Production and Operations Management
```

```
- McKinsey "AI for Social Good" (2018)
- Stanford HAI AI Index (2026)
- Microsoft "AI for Good" blog (2024)
- Salesforce.org "Nonprofit AI" case study
- Tomasev et al. (2020), Nature Communications
- Copia website
- USDA ERS food waste data (2024)
- Balcik et al. (2023), European Journal of Operational Research
- NTEN & Heller Consulting, 2024 Nonprofit Digital Investments Report
- Amnesty International report on Dutch toeslagenaffaire (2021)

Help me classify each source by type and assess its credibility. I already verified
several claims through Perplexity. Be honest about which ones are weak.
```

CLAUDE RESPONSE (SUMMARIZED)

> Here is my assessment:
>
> **High credibility (peer-reviewed, government, UN):**
>
> - Sengul Orgut & Lodree (2023) - Peer-reviewed in *Production and Operations Management*. Single food bank case study but rigorous methodology.
> - Tomasev et al. (2020) - Peer-reviewed in Nature Communications. Well-cited. Western-centric.
> - UN FAO State of Food Security (2024) - UN report. Global focus may underweight US domestic issues.
> - USDA ERS (2024) - Government data. US-only. Methodology debates exist on food waste measurement.
> - Balcik et al. (2023) - Peer-reviewed. Check if DOI is real via Perplexity.
> - Amnesty International (2021) - Reputable NGO report on documented government scandal.
>
> **Medium credibility (organizational reports, established institutions):**
>
> - Feeding America Annual Report (2025) - Self-reported but Feeding America is the largest US hunger-relief org. Incentive to show impact.
> - Stanford HAI AI Index (2026) - Academic institution. Industry-focused; less coverage of social sector.
> - McKinsey "AI for Social Good" (2018) - Consulting report. Maps AI to SDGs without measuring impact. Eight years old now.
> - Copia website - Company exists and operates, but metrics are self-reported with no third-party audit.
> - NTEN & Heller Consulting (2024) - Joint report on nonprofit digital investments. Survey of 300+ nonprofits. Good for sector-level adoption data.
>
> **Low credibility (marketing, vendor materials):**
>
> - Microsoft "AI for Good" blog (2024) - Marketing. Promoting own products. No independent evaluation of claimed outcomes.

> - Salesforce.org case study - Marketing. Selling Nonprofit Cloud. The 22% figure has no independent ROI data.
>
> I would cut your table to 10 sources for the report. Drop Balcik if you cannot verify the DOI, and drop the Amnesty report unless you use the Dutch scandal in your bias audit.

> *Good framework. I am keeping the Amnesty report because I use the Dutch scandal as an example of algorithmic harm in my bias audit. Cutting it down to 10 sources for the final table. The distinction between Feeding America (medium, self-reported but real org) vs. Salesforce (low, selling product) is useful.*

Prompt 4: SDG Strength Rating

MY PROMPT

```
Gemini mapped Second Harvest to 8 SDG targets. I think that is too many. Based on
everything we have discussed, which connections are actually backed by evidence and
which are a stretch? Rate each one as Strong, Moderate, or Weak with a one-sentence
justification. Be honest.
```

CLAUDE RESPONSE (SUMMARIZED)

> - **Target 2.1 (end hunger, ensure year-round access): Strong.** Core mission. 100M+ meals annually is direct, measurable impact on food access.
> - **Target 12.3 (halve food waste by 2030): Strong.** Food rescue diverts surplus from landfill. Most direct and measurable connection.
> - **Target 1.2 (reduce poverty in all dimensions): Moderate.** Food assistance frees household income, but the causal chain from food bank to poverty reduction is indirect. Feeding America's own research supports this but it is several steps removed.
> - **Target 17.17 (encourage effective partnerships): Moderate.** 800+ partner agencies is a genuine multi-stakeholder network. But citing partnerships as SDG progress risks circular reasoning.
> - **Target 2.4 (sustainable food production): Weak.** Food rescue is about distribution, not production systems. A stretch.
> - **Target 3.4 (reduce NCDs): Weak.** Research links food insecurity to health outcomes, but a food bank is not a health intervention. Too many causal steps.
> - **Target 4.2 (early childhood development): Weak.** Real research exists linking food security to child development, but this is not what Second Harvest directly does.
> - **Target 10.2 (social inclusion): Weak.** Plausible but generic. Any social service could claim this target.
> - **Target 13.2 (climate action): Weak.** Methane reduction from food rescue is real but incidental, not a primary operational goal.

> I recommend keeping only the 4 with Strong or Moderate ratings. Cutting 5 weak connections makes your map more credible than claiming alignment with everything.

> *This confirmed my instinct that 8 was too many. Cutting to 4: Targets 2.1, 12.3, 1.2, and 17.17. The lab says "fewer strong connections are better than many weak ones" and I agree. The SDG 17 circular reasoning point is sharp. I am keeping that connection but noting the limitation.*

STEP 6: CLAUDE (REFLECTION POLISH)

Following the lab instructions, I wrote my synthesis and reflection section entirely without AI first. Below is the prompt I used after finishing my draft to ask Claude to help me polish it.

Prompt 5: Synthesis Draft Review

MY PROMPT

```
I wrote the following synthesis paragraph for my lab report. I wrote it myself first
without AI help. Can you help me tighten it up? Do not change my voice or add any
new claims. Just help me make it clearer and check that my APA citations are
formatted correctly.

Here is my draft:

"This research showed me there is a big gap between what people say about AI in food
banks and what is actually documented. Gemini gave me a lot of sources fast,
including some peer-reviewed papers I probably would not have found on my own, but
it also presented vendor marketing right alongside academic research like they were
the same thing. It even made up a partnership between DeepMind and WFP that does not
exist. Perplexity was better for checking facts because it shows you exactly where
everything comes from, but it does not give you the big picture the way Gemini does.
Claude was most useful for the critical analysis part. The survivor bias insight,
that my research only showed food banks that succeeded with AI and completely missed
the majority that never tried it, is something I would not have gotten to on my own.

The SDG mapping turned out to be more useful than I expected, but only after I cut
it down. Gemini suggested 8 targets. After checking the evidence, I kept 4. The ones
I dropped (SDG 3, 4, 10, 13) had connections that sounded reasonable but fell apart
when you asked for proof. The framework works better as a filter than as a
checklist.

If I were advising Second Harvest today, I would say start with free tools they
already control. Google Sheets for tracking, Google Analytics for their website. The
paid platforms like Salesforce and Copia create dependency that is hard to get out
of (Iazzolino & Stremlau, 2024). Before adopting either one, they need to get data
```

```
portability in writing. What I still don't know is how they actually track
distribution now, what takes staff the most time, and whether they have tried tech
solutions before. Those questions need real conversations, not AI research."
```

CLAUDE RESPONSE (SUMMARIZED)

> Your draft is strong. The voice is clear and genuine. A few suggestions:
>
> **Tightening:**
>
> - "This research showed me there is a big gap" could be "This research revealed a significant gap" but honestly your version sounds more like a real person. Keep it.
> - "like they were the same thing" is good. Do not change this to something more formal.
> - "is something I would not have gotten to on my own" is perfect. Natural voice.
>
> **Citation checks:**
>
> - The Iazzolino & Stremlau (2024) citation is correctly formatted for APA inline.
> - You should add a citation after "Gemini suggested 8 targets" to support the SDG framework point. Consider referencing your own Gemini research or a general SDG source.
> - The Feeding America (2025) citation would fit after "100 million meals annually" if you include that detail.
> - Consider adding (Tomasev et al., 2020) after the point about AI bias to strengthen the academic grounding.
>
> **What not to change:**
>
> - Do not formalize "I would say start with free tools they already control." This directness is good.
> - Do not replace "fell apart when you asked for proof" with academic language. It is clearer this way.
>
> Overall this reads like a student who actually thought about the material, not like an AI wrote it. That is what your professor wants to see.

Glad Claude confirmed my voice is fine. I am adding the citations it suggested (Feeding America, 2025 and Tomasev et al., 2020) but keeping everything else as I wrote it. The point about not formalizing my language is a good reminder. I almost changed "fell apart" to something more academic before I sent this.

Prompt 6: Tool Comparison Paragraph

MY PROMPT

```
I also need help tightening my tool comparison paragraph. Same rules: do not change
my voice, just help me make it clearer. Here is my draft:

"Gemini was best for going wide. It found 30+ sources in minutes including academic
```

```
papers I never would have found searching manually. But it does not distinguish
between a Nature paper and a Salesforce blog post, and it fabricated at least one
claim (the DeepMind-WFP partnership). Perplexity was best for going deep on specific
claims. The inline citations made it obvious which claims had real sources behind
them and which did not. The weakness is it does not give you context or help you see
patterns across sources. Claude was best for thinking critically. The bias analysis
and corporate capture framework came out of prompting Claude with my findings and
asking it to tear them apart. The limitation is it only knows what you feed it. It
cannot go verify something on its own."
```

CLAUDE RESPONSE (SUMMARIZED)

> This is tight and well-structured already. Each tool gets one strength, one weakness, and a concrete example. Minor suggestions:
>
> - "going wide" and "going deep" is a nice parallel. Keep it.
> - "asking it to tear them apart" is vivid. Do not soften this.
> - Consider changing "It cannot go verify something on its own" to "It cannot independently verify claims" for clarity. But your version works too.
> - You might add one sentence about how the three tools worked as a pipeline rather than alternatives. Something about how each one compensated for the others' weaknesses.

Taking the suggestion to add a sentence about the pipeline. Adding: "The three tools worked best as a sequence, not as alternatives. Each one compensated for what the others missed." Keeping everything else as-is.

STEP 7: CLAUDE (FINAL POLISH)

Prompt 7: AI-Ism Check

MY PROMPT

```
I am about to submit my lab report. Before I do, scan the following paragraphs for
AI-isms. Flag anything that sounds like a chatbot wrote it rather than a student.
Look for: em dashes, overly confident declarative sentences, words like "crucial" or
"delve" or "landscape" or "underscores," short punchy sentences that all sound the
same, and any phrase I would never say out loud in conversation.

[Pasted full synthesis and reflection section, approximately 800 words]
```

CLAUDE RESPONSE (SUMMARIZED)

> I found very few AI-isms, which makes sense since you wrote the first draft yourself. A few small flags:
>
> - **"This research revealed a significant gap"** (opening sentence): "Revealed a significant gap" is borderline. "Showed me there is a big gap" was your original and sounded more natural. Consider reverting.
> - **"the documented reality"**: This phrase is fine in academic writing but check if it sounds like you.
> - **"The framework works best as a filter, not a checklist"**: This is strong and sounds like something a person would actually say. Keep it.
> - **No em dashes found.** Good.
> - **No instances of "crucial," "delve," "landscape," "underscores," "leverage," or "robust."**
> - **Sentence variety is good.** You have a mix of short and long sentences that do not follow a robotic pattern.
>
> One structural suggestion: Your reflection flows well from findings to tool comparison to recommendations. The transition from SDG mapping to tool recommendations feels slightly abrupt. Consider adding one connecting sentence.
>
> Overall verdict: This reads like a student wrote it, which is the goal. If I had to guess, I would say 90% human-written with light AI editing, which is exactly what your professor asked for.

> *Reverted "revealed a significant gap" back to "showed me there is a big gap." Added a transition sentence between the SDG mapping paragraph and the recommendations. Done. Report is ready to submit.*

Prompt 8: Citation Cross-Check

MY PROMPT

```
Final check. Here are all my inline citations and my reference list. Do they match?
Are there any inline citations without a reference, or references without an inline
citation?

Inline citations used:
- Feeding America (2025)
- UN FAO (2024)
- Sengul Orgut & Lodree (2023)
- McKinsey (2018)
- Stanford HAI (2025)
- Tomasev et al. (2020)
- USDA ERS (2024)
- Iazzolino & Stremlau (2024)
- Dhaliwal & Hou (2024)
- NTEN & Heller Consulting (2024)
```

```
Reference list:
[Pasted 10-entry reference list]
```

CLAUDE RESPONSE (SUMMARIZED)

> **Match check:** All 10 inline citations have corresponding reference list entries. All 10 reference list entries are cited inline. No orphans in either direction.
>
> **Formatting notes:**
>
> - Sengul Orgut & Lodree (2023): Correct. DOI confirmed. *Production and Operations Management* should be italicized.
> - Tomasev et al. (2020): Correct. Nature Communications volume and article number format is right.
> - Dhaliwal & Hou (2024): Confirmed as *Stanford Social Innovation Review, 23*(1), 59-60. DOI exists.
> - All other entries look properly formatted.
>
> Your references are clean. Ready to submit.

> *Fixed the Dhaliwal & Hou entry to specify Stanford Social Innovation Review as the publication. All citations match. Report formatting is done. Submitting.*