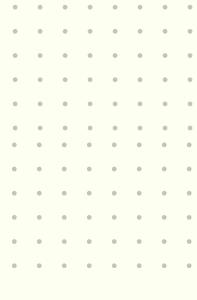


Second Bite

London's Al-powered food redistribution solution





Executive Summary

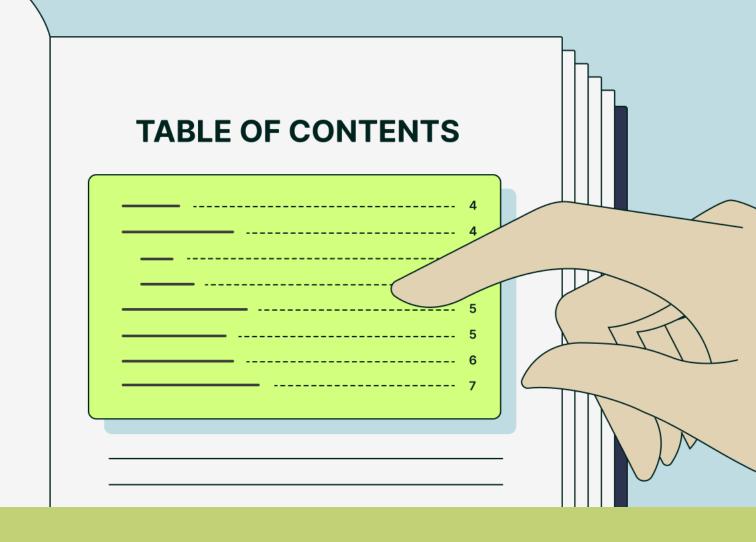
Second Bite is a socio-technical solution that harnesses AI to address the structurally intertwined challenges of food waste and food insecurity in London. Amidst an annual surplus of 1.9 million tonnes of discarded yet consumable food and a growing demographic of over two million food-insecure residents, the prevailing Surplus Food Redistribution (SFR) infrastructure is marked by fragmentation, inefficiencies, and underdeveloped digital coordination mechanisms. Second Bite operationalises a platform-based model to algorithmically mediate and optimise exchanges between commercial food donors and charitable organisations.

Central to its technical architecture is the deployment of real-time data pipelines, multivariate forecasting models (e.g., ARIMA, XGBoost), and optimisation heuristics that enable dynamic matching of supply and demand across geospatial and temporal dimensions. By integrating API-level access to donor inventory systems and employing natural language processing (NLP) to parse unstructured signals from recipient organisations, the platform constructs a continuous learning ecosystem that minimises transaction costs and enhances surplus visibility.

Informed by a normative commitment to ethical AI, Second Bite embeds explainability, data protection (GDPR-compliance), and equity-focused auditing mechanisms throughout its algorithmic governance. The platform's design anticipates scalability but also regulatory evolution, positioning itself as the foundational digital infrastructure for surplus food redistribution across the UK. In doing so, Second Bite aspires to reconceptualise surplus not as waste, but as a redistributable public good advancing a model of AI for social equity and infrastructural justice.

Al Disclaimer: for this report, ChatGPT was used to check spelling and grammar and organise notes for research.





01	Problem Formulation	03
02	Technological Solution	10
03	Organisatonal Structure	20
04	Governance Framework	24
05	Ethical Considerations	28
06	Appendix	32
07	Bibliography	40

1. Problem Formulation



Problem Formulation

1.1 Context

It is estimated that every year (WRAP, 2021) approximately 9.5 million tonnes of food are wasted annually in the UK, much of which remains fit for consumption. Concurrently, food insecurity afflicts an estimated 11-14 million individuals including 3-4 million children (Food Foundation, 2024; Butler, 2023). London offers a compelling case study of these intertwined challenges: over 2 million residents experience food poverty (London Assembly, 2022), while food waste within the capital exceeds 1.9 million tonnes annually. Surplus Food Redistribution (SFR) has been advanced as a dual-impact mechanism capable of addressing both waste reduction and food insecurity mitigation. However, despite the proliferation of redistribution organisations across the UK (see graph 1.1), surplus redistribution rates lag significantly behind those observed in the US, France, and Spain (FareShare, 2024).

Sawyerr et al. (2024) identify 3 structural constraints that reduce UK redistribution efficacy:

- Insufficient economic incentives for donors: opportunity costs of donation are significant, e.g. given alternatives like animal feed conversion.
- Inconsistent and opaque surplus visibility: unpredictability of available donations complicates operational planning for charities.
- Inefficiencies in donation decision-making among donor organisations such as supermarkets.

Again, London epitomises these systemic issues, with institutional fragmentation and underdeveloped digital coordination infrastructures hindering surplus-food redistribution.

Graph 1.1 - Current Food Redistribution Industry



1.2 Business Model

To help tackle these industry challenges, Second Bite introduces an Al-enabled, two-sided platform designed to improve surplus food redistribution by linking commercial food donors to charitable food service organisations. The platform is based on real-time and predictive surplus modelling and optimisation algorithms to dynamically match food supply with demand.

Value creation

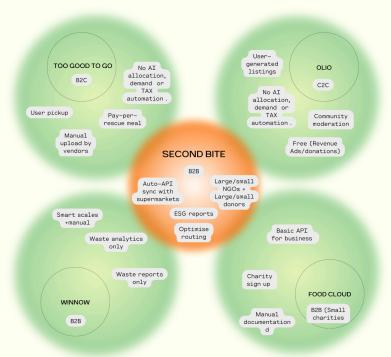
Being a two-sided platform, Second Bite is focused on generating value for both sides of the SFR market, donor organisations and charities:

Donors (Economic and Strategic Value):

- Optimised tax deductions.
- Substantial reduction in the operational overhead traditionally associated with donation workflows due to automated donor inventory data integration.
- Enhanced alignment with Environmental, Social, and Governance (ESG) objectives, strengthening public image and brand equity
- Improved compliance with emerging regulatory expectations concerning food waste and corporate sustainability

Charities (Social Value):

- Access to increased and diversified sources of food.
- Greater predictability of donations, allowing for improved planning.
- Allocation mechanisms that consider perishability, recipient capacity, and local (social) needs.



AI-Enabled Use Cases

Use Case 1 – Surplus Allocation Optimisation

- Task 1: Predict future donation volumes
- **Task 2:** Execute constrained optimisation across financial, social, and logistical objectives to optimise platform value; with reinforcement learning based on fee

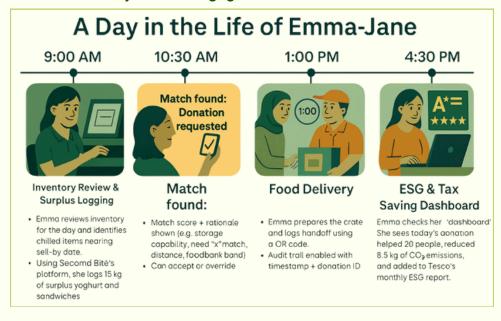
A Day in the Life of Sarah Malik 10:30 AM 11:00 AM 9:00 AM 4:30 PM 15 kg chilled yoghurt from Tesco 11:00 match score 74% Onboarding & Surplus Food Delivery **ESG Reporting Profile Setup** Notification & Verifications + Tax Incentives Logs into Second Bite Receives real-time alert Receives the delivery Learns about the benefits dashboard for the first about available food of Second Bite and and confirms its quality time and set up her surplus and taps 'Accept' various Supermarkers centre's profile and restaurants registered A Day with Second Bite on the dashboard

Graph 1.3 - Use Case Study: Allocation Optimisation

Use Case 2 - Donor Engagement Enhancement

- Task 1: Automate structured and semi-structured data ingestion through APIs
- Task 2: Provide tax savings forecasts,
- **Task 3:** Offer Al-generated, context-sensitive donation prompts, tailored to each donor's inventory and donation history, thereby nudging more frequent and impactful contributions.

Graph 1.4 - Use Case Study: Donor Engagement Enhancement



Value Capture

Value is captured from donors and charities and can be adjusted to optimise platform value:

- Commission-based revenue (ranging from 5% to 10%) on realised donor tax benefits, scaled according to donation volume and consistency.
- Tiered membership models offering access to core features.
- Premium services including surplus analytics dashboards, ESG reporting integrations, and predictive demand mapping for supply chain and CSR teams.

Growth Trajectory

Phase I – Metropolitan Scaling (London):

- Prioritise London's high-density donor and recipient networks to facilitate rapid ecosystem activation and network effect generation
- Implement donor onboarding incentives including temporary commission waivers, featured ESG recognition, and premium feature trials for large-volume early adopters

Phase II – Institutional Expansion and Policy Integration:

- Proactively align with national fiscal and environmental policy evolution, e.g. integration of potential VAT exclusions for donated food.
- Expand the platform into a unifying digital infrastructure layer across the UK's fragmented surplus redistribution landscape.

1.3 Business Architecture

Second Bite adopts a modular, Al-augmented Platform-as-a-Service (PaaS) paradigm, built to support scalable, resilient, and interoperable food redistribution. Its technical infrastructure enables real-time data flows between donors, intermediaries, logistics actors, and charitable institutions optimising coordination, transparency, and agility. Designed for extensibility, it supports flexible deployment across regions and regulatory environments. API endpoints align with standard data schemas, enabling integration with legacy and modern platforms. The architecture anticipates decentralised governance and data sovereignty challenges, ensuring long-term adaptability (Baesens et al., 2016).

Process Architecture

Supply Forecasting: Input variables include item taxonomy, volume, perishability, donor geolocation, cold-chain requirements, and donation patterns shaped by seasonality and promotions. Outputs are daily and weekly surplus forecasts by donor and item type. Models include ARIMA, Prophet, XGBoost, and anomaly detectors to flag donation anomalies (Provost & Fawcett, 2013).

Demand Forecasting: Inputs include charity data, past usage, certifications, and unstructured signals from community forums and social media, mapped to deprivation indices. Outputs consist of itemised demand vectors, urgency scores, and fulfilment likelihoods. The framework integrates probabilistic models, topic inference, and NLP to process vernacular demand signals (Snee et al., 2016).

Optimisation Engine: Optimises redistribution by aligning surplus with demand, maximising donor value (e.g., tax rebates), and minimising environmental/logistical costs. It uses reinforcement learning, combinatorial optimisation, and cost-sensitive routing, factoring perishability, traffic, and emissions (King, Keohane, & Verba, 1994).

System Outputs: The platform delivers ranked redistribution suggestions, optimised dispatch plans, and configurable dashboards with scenario simulations, early alerts, and custom analytics.

1.4 Requirements and Constraints:

Technical Requirements: The platform adheres to GDPR and ISO 27001, using encryption, federated APIs, and decentralised authentication. Modular ML workflows support continual learning. NLP and OCR process informal, multilingual input. Interfaces prioritise multilingual access, offline functionality, and visual accessibility (Myers, 2019).

Evaluation Metrics: Forecasting models are evaluated with adjusted R², RMSE, and MAPE. Classification and optimisation models are measured by precision, ROC-AUC, uplift, and equity-weighted metrics (Gerring, 2012).

Operational Requirements: Stakeholder co-design is embedded through sprints, validation workshops, and prototyping with charities, SMEs, and communities. Governance includes biannual audits, stakeholder reviews, and third-party fairness tools (Silverman, 2020).

Systemic Constraints: Challenges include fragmented data infrastructures, variable logistics capacity, and digital inequality among stakeholders—necessitating hybrid systems, interoperability solutions, and digital training initiatives (Robson & McCartan, 2016).

Ethical Governance and Mitigation Framework:

Second Bite employs layered ethical oversight. Allocation pathways are logged with metadata for auditability. Fairness monitors assess disparities across regions and demographics (Crawford, 2021).

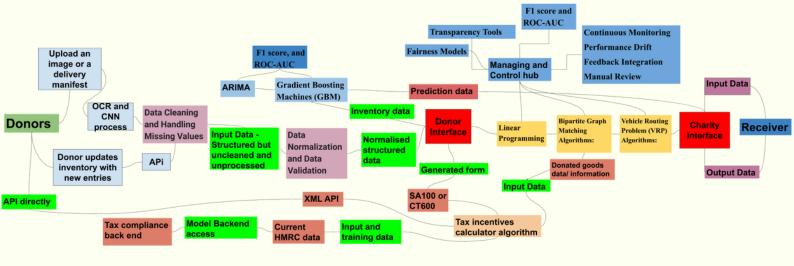
Override mechanisms enable contextual corrections by local actors. The governance model, grounded in Crawford (2021) and Eubanks (2018), prioritises justice, transparency, and participatory decision-making.

Collectively, these components form a responsible AI infrastructure, embedding equity, technical rigour, and community inclusion into all layers of the system.

2. Technological Solution



Intial Tech Overview



Graph 2.1 - Overview of the system

Design Ideology



Our design ideology is user-centricity, efficiency, and transparency. It is designed to be intuitive and accessible, from large supermarkets to small charities. Transparency builds trust and accountability among all those who use and are affected by SecondBite.







The data is delivered through either an API, Manual interface in the UI, or through uploading and image or Manifest/Excel to the techinical solution

Donor Data Types

Location Data:

Geographic-location

Inventory Data:

- Item-types
- Quantities
- Expiration-dates
- Storage-requirements
- Stock-information
- Sales-rates

Donation Data:

- Items-for-donation
- Consistent-availability
- Donation-frequency
- Use-window
- Request-data

Charity Data Types

Location Data:

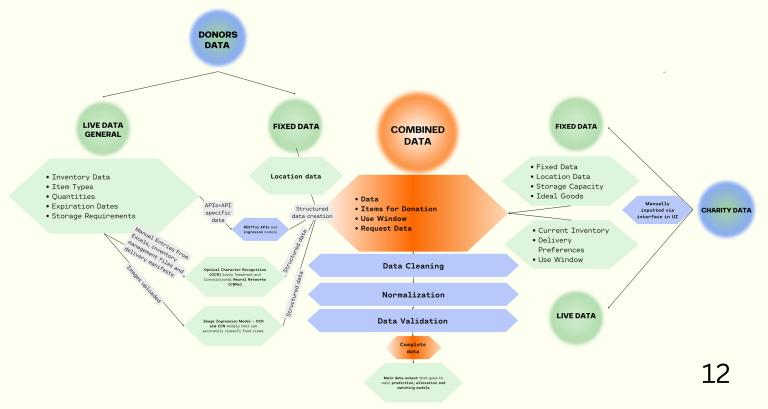
• Geographic-location

Capacity Data:

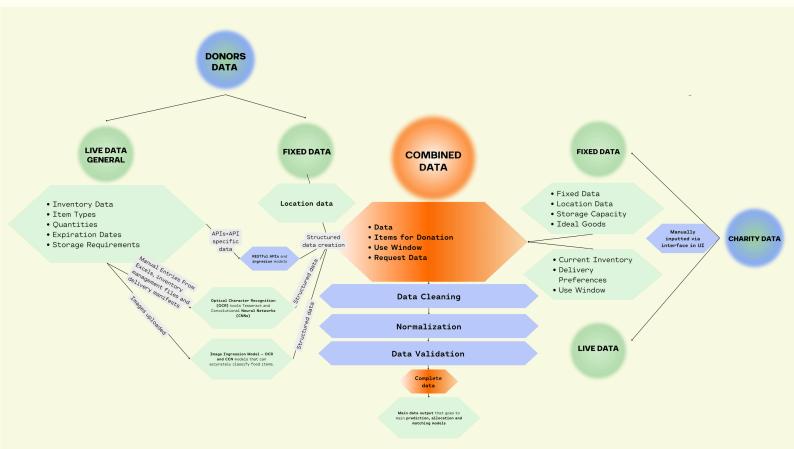
- Storage-capacity
- Current-inventory
- ideal-goods

Demand Data:

- Specific-needs
- Capacity-constraints
- Delivery-preferences
- Charity-type
- Use-window
- Consistent-demand
- Acceptance-locations



Data Sources and Data Preprocessing



Graph 2.2 - Data Sources and Data Preprocessing

Data Processing

The Data cleaning process removes duplicates and manifests missing values for data quality.

Data normalisation then standardises data for consistency in inputs.

Finally, validation approves the inputs before it reaches the models.

Outputs 🏈

There are three main branches of data outputted from the solution.

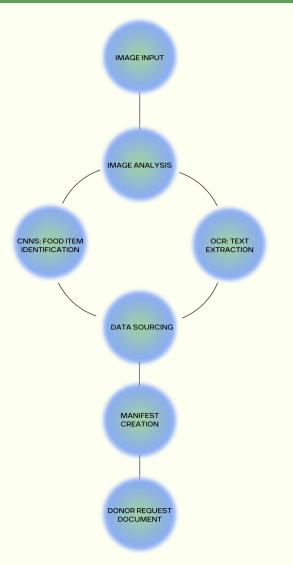
- **Forecast Data:** The donation prediction data is sent to the charity dashboard.
- Allocation Data: The matched Donors and Charities, along with delivery/pick-up information
- **Feedback Analysis Data:** Analyzable feedback data used by the evaluation system.
- Tax Data: Data on successful deliveries and tax savings generated



Overview

Our Three Ingression models function through API's or OCRS/CNNs

This ingression model can scan images of food items intended for donation and automatically generate a manifest or donor request.



Graph 2.3 - Image Input Ingression Model

The Ingression Models

These models ingest the unstructured data and transforms that data into structured formats, enabling seamless delivery to the data preprocessing.

- Form Recognition Models: OCR and CNNs process unstructured donor data like receipts and delivery manifests.
- Donor Inventory Data Ingression: Realtime API integration directly from inventory systems.
- Charity Data Ingression: Ingest the manual data entry from charities.

The ingression models function before Data Preprocessing and between the API's and interfaces from Users.



Our donor supply prediction starts with ARIMA models analysing historical data alongside inventory from the donors who use APIs from their inventory to forecast potential donation levels. Our Demand prediction employs Gradient Boosting Machines (GBM) to forecast the needs of charities based on past trends and manual charity capacity reports.

Arima

Autoregressive Integrated Moving Average models implemented using Statsmodels in Python. These models analyze historical data to predict future supply, accounting for seasonal variations and trends.

This goes to the charities on the information regarding prediction and potential supply. This is displayed on the UI of the charities interface.

Processed Data- Post structuring, cleanings and validating

Prediction algorithms

Gradient Boosting Machines (GBM):

Using XGBoost or LightGBM, these models predict demand from charities and food banks based on historical data and trends. GBMs handle complex relationships and interactions between features, providing robust and accurate forecasts. They are well—suited for scenarios where the demand patterns are influenced by multiple factors.

FI score, and ROC-AUC

Used to monitor and improve model performance. These metrics help in assessing the effectiveness of the models and ensuring they meet the desired accuracy and reliability standards.

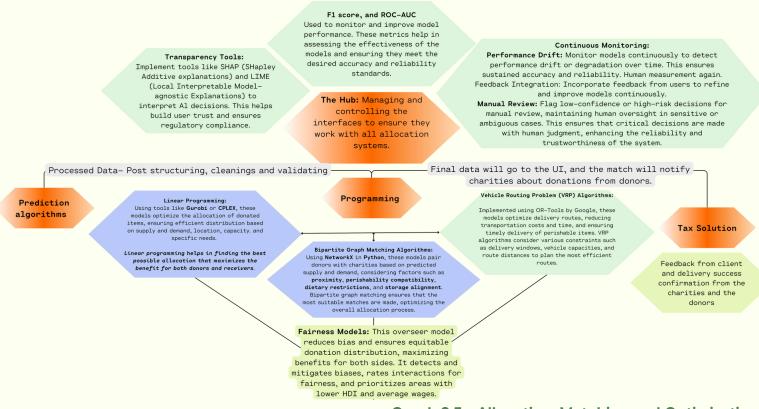
Graph 2.4 - Prediction Solution

These models inform donors when to donate and when charities should have delivery or pickup staff on call. In the future, we could alter the rate we take on tax savings to incentivise donors to donate on low supply periods.





Optimises the distribution of donated items to charities based on various factors.



Graph 2.5 - Allocation, Matching and Optimisation

Model Design



Allocation optimization initally uses linear programming like Gurboi to prime the donors and charities profiles using the supply, demand, location, and capacity data.

The Bipartite Graph Matching to assign the charities to the donors from this data.

Then the Route optimization with VRP algorithms plans the most efficient delivery routes and recomeends pick up times.

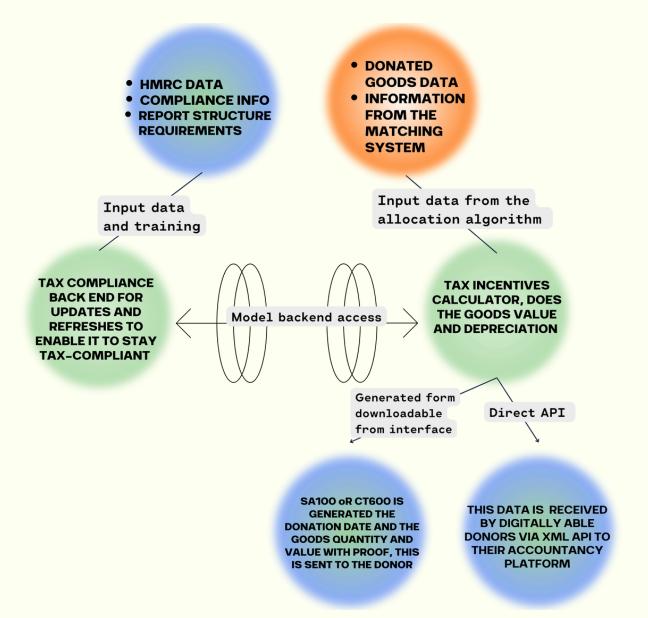
Optimisation Metrics

What is the goal of the software, how do we measure success for our models. We tier our three success metircs as follows:

- 1. Tax Value/quantity of food delivered
- 2. Distance and ease of delivery/pick up
- 3. Existing relationships with successful deliveries in the past.



Automates tax receipt generation and calculates tax incentives for donations.



Graph 2.6 - Tax Solution

In the UK, businesses can deduct the cost of donated trading stock from their total business profits before paying tax, reducing their taxable profit. UK Government (2024). Our tax solution assigns a unit cost value to each donated good and provides this value as a tax deduction. Remaining life is based on the food stuff's usability from its stock or expiry date depending on the available data. This donation has to be to a registered charity and recorded in detail, both things our solution accomplishes.

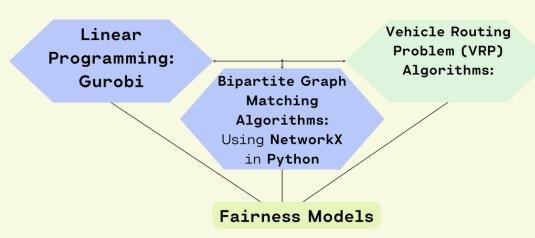
We support both generations of Self Assessment Tax Return (SA100) for sole operators and Corporation Tax Return (CT600) forms for larger firms. For digitally-able firms, we connect directly to accountancy systems like Xero via XML API to deliver their deductions directly.



Depreciated Value = Σ (Unit Cost × Quantity Donated)



Evalution, Fairness and Management Solution



Graph 2.7 Fairness mode interactions

Fairness Model 🌉

This overseer model actively aims to mitigate bias and maintain equitable donations. It rates each match for fairness and prioritizes areas with lower HDI, fewer donors, and average wages with more donations.

Transparency Tools:

Implement tools like SHAP (SHapley Additive explanations) and LIME (Local Interpretable Model-agnostic Explanations) to interpret Al decisions. This helps build user trust and ensures regulatory compliance.

F1 score, and ROC-AUC

Used to monitor and improve model performance. These metrics help in assessing the effectiveness of the models and ensuring they meet the desired accuracy and reliability standards.

The Hub: Managing and controlling the interfaces to ensure they work with all allocation systems.

Continuous Monitoring:

Performance Drift: Monitor models continuously to detect performance drift or degradation over time. This ensures sustained accuracy and reliability. Human measurement again. Feedback Integration: Incorporate feedback from users to refine and improve models continuously.

Manual Review: Flag low-confidence or high-risk decisions for manual review, maintaining human oversight in sensitive or ambiguous cases. This ensures that critical decisions are made with human judgment, enhancing the reliability and trustworthiness of the system.

Graph 2.8 The Hub and its compotents

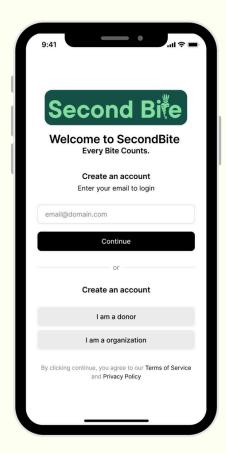
Evaluation Hub (

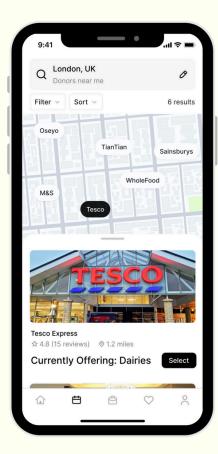


The evaluation solution receives KPIs and metrics such as accuracy, precision, recall, F1 score, and ROC-AUC to measure model effectiveness. Czakon, J. (2024). There are regular bias and fairness audits and tools that allow human-in-the-loop interventions to stop discriminatory results. The continuous monitoring and feedback help refine model performance and prevent performance drifting.



Ul Functionality and Interface Design





Charity Side of the UI

Sign-Up and Profile Management

- Registration Form
- Profile Dashboard

Capacity and Demand Management

• Data Entry Form: storage capacity, current inventory.

Donation Requests

- Request Form:
- Request History
- Request Confirmation and Notifications
- Map of Donors

Feedback and Communication

- Feedback Form
- Messaging System



Donor Side of the UI

Sign-Up and Profile Management

- Profile Dashboard details
- Overview of donation history



Inventory Management:

- Data Entry Form: Manual donation entries
- Upload Functionality: CSV or Excel files.
- Image Upload.
- API Integration: Interface to connect to via RESTful APIs.

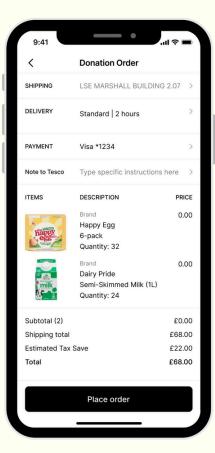
Donation Management.

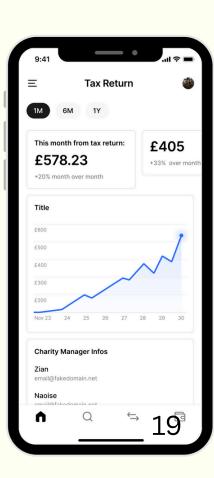
- Donation History
- Donation Confirmation and Notifications
- Map of charity locations for delivery and Pick up

Feedback and Communication

- Feedback Form
- Messaging System







3. Organisational Solution





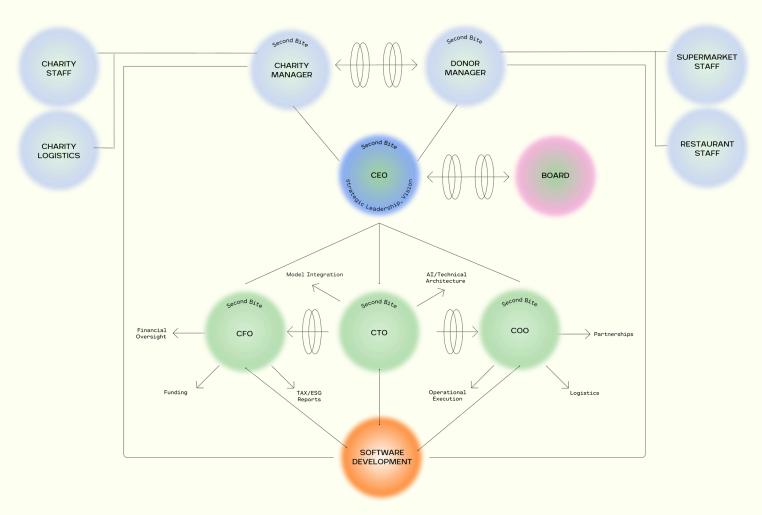
3.1 Second Bite Organisation

The Second Bite organisation consist of business and technology functions (see chart 3.1).

Second Bite is a social enterprise aimed at helping to solve social problems with the aid of technological innovation.

To ensure technology is in service of the social-entrepreneurial goals, product development will be driven by business functions (Donors & Charities Managers) in cooperation with technical design leads.

Graph 3.1 - Organisation Chart



3.2 Organisational Impact on Stakeholders

As can be seen from the overview of Second Bite's Work Systems (Appendix 2), the Second Bite Al systems interact with donors and charities extensively, affecting their work systems ((compared to current SFR industry practices) primarily across 3 areas: **Donation Data Collection**, **Donation Recommendations**, and **Delivery**.

Donation Data Collection

Current SFR industry practices require donors to manually indicate what produce they wish to donate.

As graph 3.2 shows, Smart Bite offers 2 additional options for donation data collection to significantly reduce resource requirements:

- Use APIs to automatically provide inventory data to the Second Bite platform.
- Send images of produce to be donated through the mobile interface.

These additional options can significantly reduce resourcing requirements for the donation process.

Donation Recommendations

Current SFR industry practice sees the utilisation of manually-inserted data on produce to be donated/

Second Bite's ability to extract data automatically from inventory systems enables the AI systems to proactively make recommendations to the donor as to what inventory to donate (see graph 3.3).

Current SFR industry practice Manual provision of data on produce to be donated **Second Bite Optimisation** Manual provision of data on produce to be donated Images of produce to be donated Automated extraction

Graph 3.2 - Organisational Impact Data
Collection

Current SFR industry practice

Model allocates produce offered



Second Bite Optimisation

Models allocate produce offered (for manual or image input)

Models make optimised recommendations based on inventory data (for APIs)





Especially where the donor has set no/limited donation parameters, Second Bite AI systems have more autonomy to make optimal donation recommendations (see Appendix 3 for further details). This can optimise tax benefits but also impact donor inventory management processes and indeed, create role-boundary changes as individuals in donor organisations must interact with information produced by AI systems to make financial optimisation decisions ("donate or sell"). Donors making use of data extraction will receive training on how to engage effectively with the AI systems.

For charities, recommendations are not expected to yield organisational changes compared to current industry practices.

Delivery

As part of the optimisation process, Second Bite considers optimal means and routes of delivery, providing stakeholders with donation and transport recommendations that reduces resourcing spent on delivery or collection of donated food.

4. Governance Framework and Legal Considerations



4. Governance and Legal Framework

Falling under the purview of our Board, CTO, and COO, legal standards compliance, protecting stakeholders, and maintaining the operational ethics of our system are of utmost importance. As we develop, our CTO will cover functional modifications and our COO will implement policies for error handling, complaints procedures, and regulatory compliance, particularly concerning data governance and AI ethics. All are under the supervision of our board.

- **4.1** Functional updates and changes will follow structured governance models ensuring compliance. Algorithmic improvements, data processing enhancements, or new reporting features undergo thorough impact assessments. All changes will be tested in controlled environments before deployment and stakeholders will be notified ahead of major updates. Quarterly Al bias audits will identify and rectify any unintended disparities in food allocation.
- **4.2** Since our platform processes sensitive data, including donor records, food donation logs, and route information, we will adhere to General Data Protection Regulation (GDPR) and NIST AI RMF principles, and make any adjustments to account for the upcoming 2025 AI Act. Data minimization collects only necessary information, and all data storage is encrypted within the Google Cloud Platform (GCP) meeting ISO 27001. Users will have the right to access, correct, and request deletion of their data through Zendesk, ensuring transparency and privacy protection (Infrastructure Framework Appendix 4.1)

Additionally, our system will align with the UK Food Safety Act (1990) and Food Standards Agency (FSA) guidelines, ensuring donated food meets food safety standards. Al flags items nearing expiration preventing unsafe food from being redistributed. Traceability measures to track donation destinations and food sources comply with the DEFRA Food Recovery Hierarchy (Beckmann et. al, 2021).

Finally, tax and legal integration will be streamlined through an HM Revenue and Customs (HMRC) compliant automated tax receipt generator. Donations will be digitally documented, and tax incentives will be automatically calculated. Our API will connect directly to Charities Online XML to facilitate tax claims, ensuring regulatory alignment.

REGULATIONS

IMPLICATIONS FOR SECOND BITE

EU AI ACT

Second Bite's Al-driven allocation system falls under the EU's high-risk classification due to its impact on welfare distribution. This requires: (1) Explainable AI (e.g., providing NGOs with clear reasons for each match, such as proximity or dietary alignment), (2) Human-in-the-loop validation (NGO staff must confirm/reject AI proposals, with rejections triggering model retraining), and (3) Annual bias audits to ensure no demographic disparities. Non-compliance risks fines up to €10M or 2% of global revenue.

JK GDPR

Limits how Second Bite collects and processes data for demand forecasting. NLP models must use aggregated, anonymized data from social media or public forums, excluding personal identifiers. Beneficiary records stored by partner NGOs must be deleted after 6 months unless explicit consent is obtained. Violations can result in fines of up to £17.5M or 4% of annual revenue.

SAFETY ACT (1990)

Mandates strict control over food handling and storage. Second Bite must: (1) Monitor real-time temperature logs (±2°C for chilled items) during transit, (2) Document handover protocols to charities, and (3) Ensure no spoiled food reaches recipients. Liability extends to donated goods; NGOs can litigate if illnesses occur. Compliance adds ~8% to logistics costs but prevents legal penalties.

ENVIRONMENTA L PROTECTION ACT + WRAP

Requires Second Bite to quantify and report diverted food waste using WRAP's GHG emission metrics (e.g., "X tonnes CO_2 e saved"). The UK's 2030 target of 50% waste reduction necessitates dynamic routing AI to minimize spoilage. Alignment unlocks access to government grants and enhances ESG branding.

CHARITIES ACT

Ensures donations serve charitable purposes (e.g., poverty alleviation). Second Bite must: (1) Track calorie-per-recipient ratios, (2) Avoid allocations that could harm health (e.g., sugary foods to diabetic communities), and (3) Disclose impact metrics (meals delivered/ $\mathfrak L$ spent). Violations risk loss of charitable status for partners.

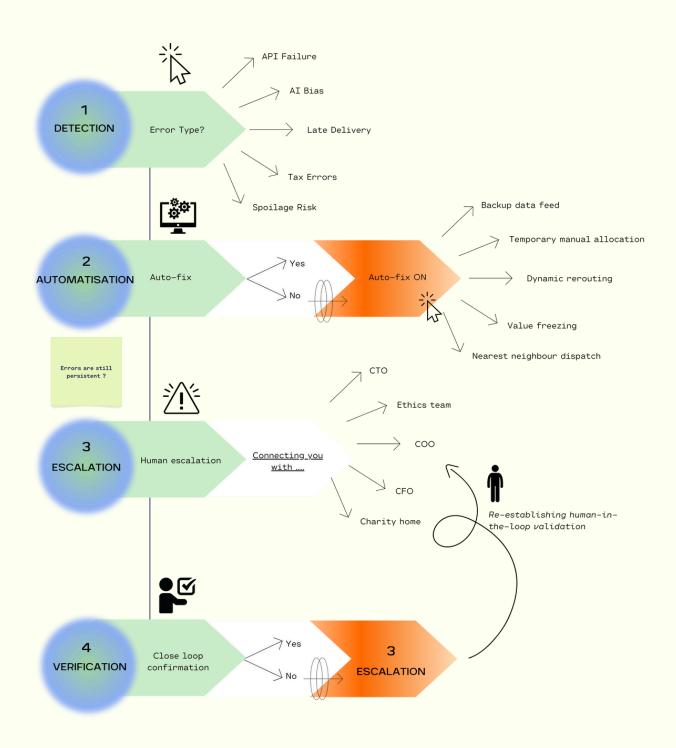
HMRC TAX COMPLIANCE Donated food must be valued accurately for tax deductions. Second Bite's Al calculates depreciation based on: (1) Expiry timelines (0% depreciation if >1 week fresh), (2) Nutritional value (higher deductions for staple foods), and (3) Market alternatives. Discrepancies >10% trigger audits. Automated reports reduce donor paperwork.

ISO 27001 (Voluntary)

While optional, certification strengthens partnerships with supermarkets by ensuring end-to-end data encryption and secure API integrations. Requires annual penetration testing and audit trails. Adds 5–10% to IT costs but is critical for scaling into EU markets.

4.3 Error handling

An error-handling framework maintains the reliability of SecondBite's AI system as errors like incorrect allocation have significant consequences (Amodei et al., 2016). Our system will employ real-time monitoring, automated anomaly detection, and human oversight for critical decision-making. If errors occur, we will implement a rapid-response protocol that includes logging, analyzing, and correcting the issue, followed by an audit to prevent recurrence.



Graph 4. 2- Error handling

5. Social and Ethical Considerations



5. Ethical & Social Considerations

While AI and data-driven infrastructures offer promise in optimising redistribution, they are embedded within power-laden knowledge regimes. As Eubanks (2018) and Seaver (2017) argue, algorithmic systems are not value-neutral; they reflect and potentially entrench systemic biases by learning from skewed datasets. In light of these dynamics, these are the ethical challenges that Second Bite ought to tackle to ensure fairness:

5.1 Bias

Our algorithms learn from historical donation patterns to match donors to charities. However, if donations have historically taken place between large supermarkets and charities, these parties could become over-represented in our datasets. Moreover, representation bias could be exacerbated by our data collection methods: Second Bite gathers structured data from supermarkets via APIs, whereas smaller vendors have to rely on manual input. This inconsistency may exclude smaller donors from reaching food-insecure communities. Training our models on biased data could create feedback loops that further exclude weaker actors from donation networks.

5.2. Accountability

Our system consists of an ecosystem of actors in which accountability may become diffused. The technical structure of Second Bite, which contains multiple layers of algorithmic models for supply/demand prediction, allocation, and route optimisation, introduces many decision points. If a delivery fails to arrive, or if a high-need charity is excluded, it may become unclear whether the error arose from incorrect data input, faulty prediction models, or false allocation. Without explainable AI (XAI) mechanisms in place, stakeholders may struggle to understand how or where errors have occurred, and become unable to contest them. Such an accountability gap could erode stakeholder trust in our platform and harm vulnerable communities that are in dire need of surplus food.

5.3 Transparency

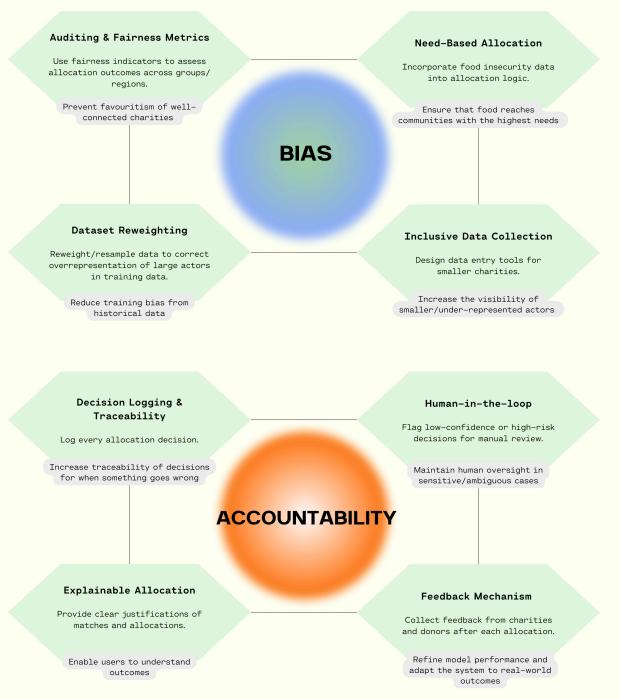
Second Bite's algorithmic decisions influence the allocation of surplus food to charities and NGOs. Without sufficient transparency, these opaque decisions become inexplainable to actors with limited technological expertise or power within the ecosystem (Ananny & Crawford, 2018). A lack of clarity may undermine trust, discourage meaningful participation, and widen the power gap between donors and recipients (Eubanks, 2018). Moreover, opacity may concentrate power in the hands of developers or platform operators and remove those who are most heavily affected by algorithmic decisions from the decision-making process. Therefore, transparency is not only a functional requirement but also a social obligation (Pasquale, 2015). Ultimately, transparency goes beyond mere technical explainability, and requires making Al-powered decisions that all stakeholders may understand.

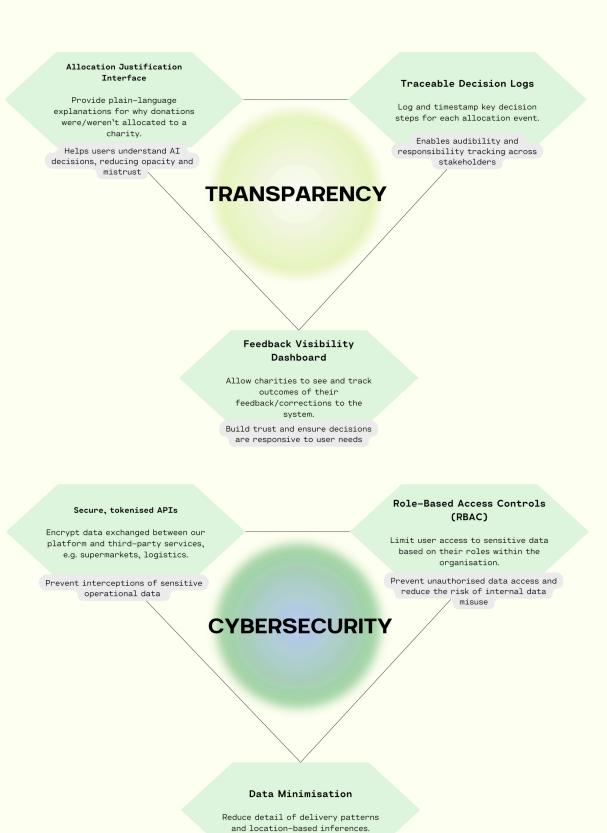
5.4 Cybersecurity

While Second Bite utilises encrypted data, secure APIs, and rigorous access controls, security risks that could affect donors and charities remain. Firstly, our platform's integration with third-party logistics and inventory management systems may create entry points for unwanted access to or exploitation of our donors' sensitive data. Insider misuse poses another threat, where authorised personnel may access or manipulate confidential operational data. Furthermore, another concern is metadata exposure, where anonymised data – such as location and delivery frequency – may reveal sensitive data about vulnerable communities and charities. If leaked or misused, this information could put our users at risk of profiling or stigmatization.

5.5 Solutions

The aforementioned social and ethical concerns underscore the importance of creating rigorous ethical guidelines that safeguard the rights and interests of all the stakeholders in our ecosystem. We propose the following technical solutions to address these challenges:





Graph 5.1 - Navigating Ethical Challenges

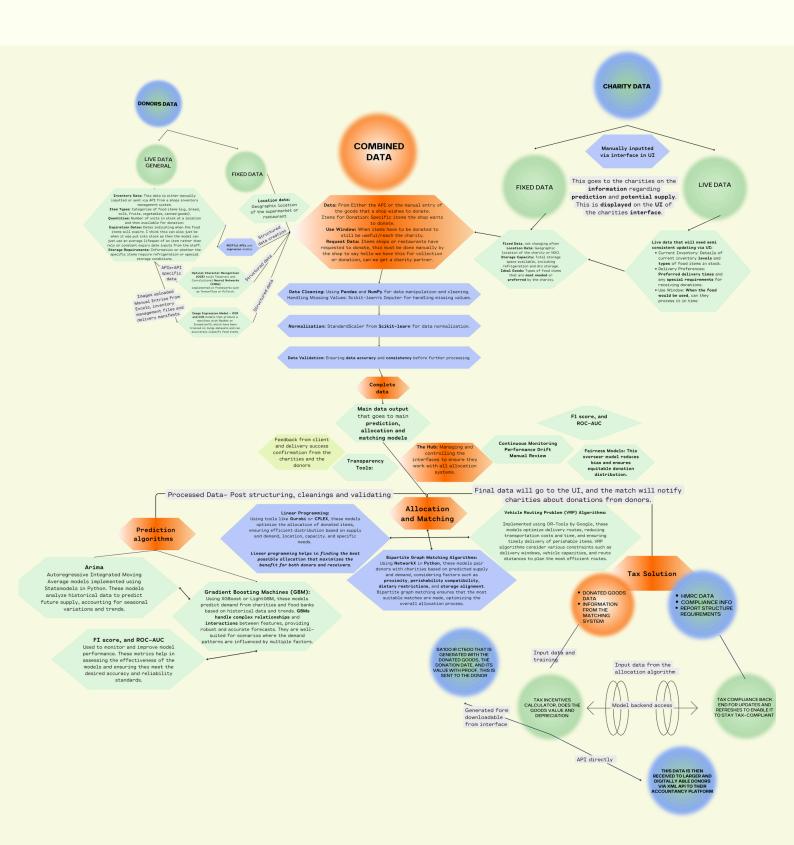
In sum, algorithmic fairness must be complemented by fair datasets, accountability, contestability, and traceability. Ultimately, Second Bite aspires not only to rectify inefficiencies in food redistribution, but also to embody a model of ethically grounded AI infrastructure.

Prevent re-identification of vulnerable communities and protect them from profiling

Appendices



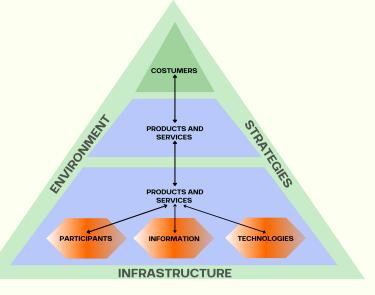
Appendix 1 - Complete Technological Overview



Appendix 2 - Second Bite Work Systems Framework

The Work Systems Framework (Alter, 2013) was designed to display an organisation's work systems, including (relationships between) customers & products, major processes and activities, as well as participants, information and technology.

Below this Framework is applied to the Second Bite platform:



Second Bite Work Systems Framework

Customers

- Individuals facing food insecurity
- Charities
- Donor organisations

Products/Services

- (Automated) food surplus identification
- Optimisation of allocation of surplus food to charities
- Tax deduction optimisation for donors

Major Processes and Activities

- Donor provides food surplus manually (structured or images)
- 2nd Bite API extracts data from donor inventory mgmt. system
- Charity provides demand data
- 2ndBite algorithm predicts supply and demand (p/day, p/product cat)
- 2ndBite algorithm optimises allocation (based on platform value)
- 2ndBite algorithm makes suggestions to donor to further optimise allocation
- Donors confirm, adjust or reject donation recommendations.
- Charities confirm or reject allocation.

Participants

- Donation managers (donor orgs)
- Supplies managers (charities)
- Second Bite employees

Information

- Inventory data
- Charities' food needs data
- Donation supply predictions
- Donation recommendations
- Feedback data (rejections of recommendations)

Technology

- Inventory mgmt. APIs
- Image recognition model
- Supply prediction model
- Allocation optimisation algorithm
- Cloud & Internet

Appendix 3- Assessing Impact on Work Systems

To assess the organisational and role-boundary impact of Second Bite's Al systems, the 'smartness of Al' concept - developed by Alter (2022) was deployed.

This concept scores the degree to which AI systems can operate autonomously - as opposed to following (human) instructions - across **4 core capabilities**:

- Information processing: extracting, manipulating information
- **Self regulation:** self-correction of errors
- Knowledge acquisition: inferring from data
- Action in the world: e.g. making recommendations.

To assess organisational impact of Second Bite's Al systems on stakeholders, key interaction points between stakeholders (donors and charities) and Second Bite Al systems were each scored for these 4 capabilities (see tables below).

The greater the autonomy of the AI system, the less it follows stakeholder instructions, and thus the greater the likelihood for it to impact on existing work systems and roles.

As can be seen from Table 4, AI system autonomy is significant where donor inventory data is extracted automatically through APIs and recommendations about donation are made to donors without pre-set constraints/parameters. This autonomy can impact inventory-management processes of donors, as Second Bite AI systems 'peek across the fence' and make inventory management recommendations. This can also cause role boundaries for individuals in donor organisations to change. Therefore, training on how to interact with the AI systems is offered to donor organisations that make use of APIs to provide inventory data (especially those that do not set donation parameters).

Table 1: Donors - Manual inputting of data on produce to be donated

prescriptions Does not perform			
Follows prescriptions		,	
Uses prespecified			
Partially scripted			
Unscripted			

Algorithmic Mgmt

> Human Mgmt

Table 2: Donors - Providing images of produce to be donated

Unscripted				
Partially scripted				
Uses prespecified inputs				
Follows prescriptions				
Does not perform				
	Self Regulation	Information Processing	Knowledge Acquisition	Action in the world

Algorithmic Mgmt

> Human Mgmt

Table 3: Donors - Automated inventory data extraciton (APIs) with donation parameters set by donor

Unscripted					Algorithmic
Partially scripted					Mgmt
Uses prespecified inputs					
Follows prescriptions					
Does not perform					Human
	Self Regulation	Information Processing	Knowledge Acquisition	Action in the world	Mgmt

Table 4: Donors - Automated inventory data extraction (APIs) without donation parameters set by donor

Partially scripted Uses prespecified					Algorithmic Mgmt
Follows prescriptions					
Does not perform					Human Mgmt
	Self Regulation	Information Processing	Knowledge Acquisition	Action in the world	27

Table 5: Charities - Allocation of donations

Unscripted					
Partially scripted					Algorithmic Mgmt
Uses prespecified inputs					
Follows prescriptions					
Does not perform					Human Mgmt
	Self Regulation	Information Processing	Knowledge Acquisition	Action in the world	

Appendix 4 - Infrastructure Framework

FUNCTION	TECHNOLOGY
CLOUD COMPUTING	Google cloud computing (GCP)
DATA MANAGEMENT	GCP (SQL) Cloud Firestore (NoSQL)
OPERATIONS (LOGISTICS AND TRANSPORT)	Firebase
CYBERSECURITY	SentinelOne
TAX AND SECURITY	Thomson Reuters
COSTUMER SERVICE	Zendesk
API CONNECTIONS	HMRC Agent Authorization API Charities Online XML API

Bibliography



Alter, S.(2013) "Work System Theory: Overview of Core Concepts, Extensions, and Challenges for the Future," Journal of the Association for Information Systems, 14(2), .DOI: 10.17705/1jais.00323

Alter, S. (2022) "Understanding artificial intelligence in the context of usage: Contributions and smartness of algorithmic capabilities in work systems" International Journal of Information Management 67(2022) DOI: https://doi.org/10.1016/j.ijinfomgt.2021.102392

Baesens, B. et al. (2016) 'Transformational Issues of Big Data and Analytics in Networked Business', MIS Quarterly, 40(4), pp. 807–818. https://www.jstor.org/stable/26629677

Beckmann, J., Shapiro, M., Angel, S., Eigen, M., & Broad, E. (2021). UNITED KINGDOM Butler, P. (2023), "Number of UK children in food poverty nearly doubles in a year to 4m", The Guardian, 1 March, pp. 2-4

Crawford, K. (2021) The Atlas of Al: Power, Politics, and the Planetary Costs of Artificial Intelligence, JSTOR. Yale University Press. https://www.jstor.org/stable/j.ctv1ghv45t

Czakon, J. (2024) F1 Score vs ROC AUC vs Accuracy vs PR AUC: Which Evaluation Metric Should You Choose? Neptune.ai. https://neptune.ai/blog/f1-score-accuracy-roc-auc-pr-auc

FareShare (2024) "Where's the Food? Strengthening Communities through Surplus Food Redistribution" available from: https://fareshare.org.uk/wp-content/uploads/2024/06/Wheres-the-Food-FareShare-Manifesto-190724.pdf

Fawcett, T. and Provost, F. (2013) Data Science for Business, ResearchGate. O'Reilly Media Inc. https://www.researchgate.net/publication/256438799_Data_Science_for_Business

Gordon, F. (2019) 'Virginia Eubanks (2018) Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. New York: Picador, St Martin's Press', Law, Technology and Humans, 1(1), pp. 162–164. https://doi.org/10.5204/lthj.v1i0.1386

King, G., Keohane, R.O. and Verba, S. (1994) Designing social inquiry: Scientific inference in qualitative research. Princeton, NJ: Princeton University Press.

Myers, M.D. (2013) Qualitative Research in Business and Management. SAGE.

Pasquale, Frank. The Black Box Society, 5 Jan. 2015, https://doi.org/10.4159/harvard.9780674736061

Sawyerr, E., Bourlakis, M., Conrad, D. and Wagstaff, C. (2024) "Impact pathways: unravelling the hybrid food supply chain – identifying the relationships and processes" International Journal of Operations & Production Management 44(7) DOI: https://doi.org/10.1108/IJOPM-05-2023-0362

Seaver, N. (2017) 'Algorithms as culture: Some tactics for the ethnography of algorithmic systems', Big Data & Society, 4(2), pp. 1–12. https://doi.org/10.1177/2053951717738104

Snee, H. et al. (eds) (2016) Digital Methods for Social Science. London: Palgrave Macmillan UK. https://doi.org/10.1057/9781137453662

UK Government (2024). Tax when your limited company gives to charity.available from: https://www.gov.uk/tax-limited-company-gives-to-charity