Integrated Ontological Framework for Food Waste Prevention, Recovery, and Valorization

Emily Steliotes[1,2], Riley Lawson[1] and Matthew Lange[1,*]

[1]International Center for Food Ontology Operability Data and Semantics (IC-FOODS), Davis, CA, USA
[2]Department of Food Science and Technology, University of California-Davis, Davis, CA, USA

Abstract

Food loss and waste (FLW) is a global sustainability challenge spanning the food supply chain from farm to fork. We present an integrated ontological framework that defines and interrelates top-level classes for food waste generation, prevention, recovery, and valorization within a Basic Formal Ontology (BFO) alignment. By reusing and extending Open Biological and Biomedical Ontology (OBO) Foundry ontologies—including FoodOn for food materials, ENVO for environmental/waste terms, IAO for information artifacts, and RO for standard relations—as well as sustainability-focused ontologies (SCO, SDGIO, SuMSO, SDC), the framework ensures semantic interoperability across domains. This ontology addresses known gaps in existing FLW ontologies (e.g. lack of prevention/recovery concepts, limited process modeling, poor cross-domain term alignment). We incorporate domain knowledge from the Upcycled Food Standard (2025), FAO, USDA, and the Food Loss & Waste Accounting Standard to ground the ontology in real-world definitions and reporting practices. The result is a high-level ontology that supports cross-sector traceability and sustainability reporting, using key object properties like hasWasteStream, derivedFrom, and eligibleForCertification to model waste flows, derivation of upcycled products, and certification eligibility. We show how this semantic framework can support AI, large language models (LLMs), and linked-data applications by providing a common, machine-readable representation of food waste concepts. An excerpt of the ontology's Turtle (TTL) code is provided to illustrate the structure.

Keywords

food loss and waste, food ontologies, waste ontologies

1. Introduction

Food loss and waste (FLW) represents a significant inefficiency in global food systems, with roughly one-third of food produced never being consumed. This inefficiency not only implies lost nutrition and economic value, but also exacerbates environmental impacts through wasted resources and greenhouse gas emissions [1]. International targets like United Nations SDG 12.3 call for halving per-capita global food waste by 2030, prompting governments and organizations worldwide to set FLW reduction goals.

Achieving these goals requires coordinated efforts to measure, track, and intervene in FLW across all stages of the supply chain. However, a persistent barrier is the lack of a common ontology—a shared semantic framework—to integrate data and knowledge about how, where, and why food is wasted. Without standardized terms and relationships, data on wasted food remains siloed among farmers, processors, retailers, waste managers, and policymakers, undermining collective action.

Efforts to formalize knowledge about FLW in ontologies have begun, but recent reviews highlight critical gaps. Lange et al. find that only a handful of ontologies (eight identified in their review)

16th International Conference on Biological and Biomedical Ontology (ICBO 2025)—Food, Waste, and Sustainability: Synergizing Ontology Efforts 2025 Workshop, November 14, 2025. Virtual conference hosted at the PTH Brixen College, Bressanone, Italy.

*Corresponding author.

emily@ic-foods.org (E. Steliotes); riley@ic-foods.org (R. Lawson); matthew@ic-foods.org (M. Lange)

0009-0006-0501-880X (E. Steliotes); 0009-0007-8463-1772 (R. Lawson); 0000-0002-6148-7962 (M. Lange)

specifically address food waste, and most focus narrowly on waste *valorization* (e.g. turning waste into bio-products) [2, 3]. Few ontologies explicitly support waste prevention or reduction, and none model food recovery/rescue—even though prevention and recovery are the top-priority strategies for mitigating FLW. This misalignment with the accepted food waste hierarchy (which prioritizes preventing waste at the source, then rescuing edible surplus, and lastly recycling/valorizing residuals) reveals a major ontological gap [1]. For example, the US EPA's "Wasted Food" hierarchy emphasizes prevention, donation (recovery), and upcycling as the most preferred pathways, yet existing ontologies underrepresent these upstream interventions [1, 2]. In addition, current FLW ontologies tend to emphasize static categorizations of waste (e.g. source type or composition) and overlook process-based dynamics such as where waste goes, when it occurs, and what causes it. These dynamic descriptors—including waste flow destination, timing, environmental impact, and causal antecedents—remain almost entirely unaddressed in prior models, resulting in only partial traceability of waste streams and limited utility for informing prevention or cross-sector solutions.

Another shortcoming is poor interoperability and misaligned terminology. Many earlier efforts were developed in isolation, without adhering to established ontology frameworks or integrating with related domains. As a result, important concepts like "food", "waste", or "sustainability" are defined inconsistently, hampering data integration across stakeholder domains. For instance, FLW ontologies have shown limited alignment with OBO Foundry ontologies, meaning they do not share a common upper ontology or utilize the Relations Ontology (RO) for standardized relations [2, 3]. They also often lack mappings to external vocabularies (e.g. agricultural, environmental, or supply chain ontologies), making it difficult to link food waste data to broader food system knowledge. In practice, the absence of a shared semantic foundation means that a "food waste" dataset in one system may not be readily comparable or connectable to data in another (e.g. linking a retailer's waste records to a food rescue organization's data). This challenge extends to differing definitions of FLW: for example, FAO defines "food loss" as losses that occur from harvest up to (but not including) retail, and "food waste" as losses at the retail and consumer end [4, 5]. Meanwhile, US agencies sometimes use "food waste" as a blanket term or prefer "wasted food" to emphasize residual value [6-8]. A robust ontology must reconcile these nuances by providing clear, reference-backed class definitions for the concepts of food loss and food waste, ensuring all stakeholders "speak the same language".

To address these gaps, we propose an Integrated Ontological Framework for Food Waste Prevention, Recovery, and Valorization. This framework is designed as an upper-level domain ontology that aligns with the Basic Formal Ontology (BFO), a widely used upper-level ontology, to facilitate rigorous classification of both material entities (like food waste itself) and processes (events and activities like generation or recovery) [6] The ontology builds upon existing OBO Foundry ontologies: we reuse classes from FoodOn (for food materials/products), ENVO (for waste materials and environmental contexts), the Relations Ontology (RO) (for standard relationships such as part-whole and derivation), and IAO (for information entities, e.g. data records or standards) [7-10].. In addition, we extend concepts from specialized sustainability ontologies: the Sustainability Core Ontology (SCO) and domain ontologies like SDGIO (which provides terms for Sustainable Development Goals, including targets on food waste), Sustainable Meat Systems Ontology (SuMSO), and the Sustainable Development and Climate (SDC) Ontology [11-14]. By integrating these, we ensure that our framework can interface with sustainability metrics (e.g. linking a food waste class to SDG 12.3 indicators) and sector-specific contexts (e.g. SuMSO concepts for livestock supply chain waste).

Our ontology explicitly covers the full spectrum of FLW mitigation activities: from the generation of food waste (as an outcome of other processes) to interventions for preventing waste, recovering surplus food for human use, and valorizing waste into new products or resources. We define top-level classes for each of these and formally relate them (e.g. linking a *Food Waste Generation* process to the *Food Waste* material it produces, or linking a *Food Waste Recovery* process to the waste input it redirects). Importantly, the ontology encodes dynamic aspects like waste destinations (e.g. landfills, animal feed, donation) and causal factors (e.g. spoilage due to lack of refrigeration) through relationships and contextual classes. This approach answers recent calls to move beyond cataloging waste "as it exists" towards modeling waste "as it happens"—capturing flows, destinations, and drivers in a machine-readable, interoperable manner. By adhering to OBO Foundry best practices and linking to external ontologies, our framework aims to provide the "structured data ecosystem" needed for cross-sector interoperability in FLW data.

Finally, we note that a robust FLW ontology has implications for emerging technologies in data analysis and knowledge management. Large language models (LLMs) and AI systems stand to benefit from a well-structured ontology: such an ontology can serve as a knowledge graph backbone that grounds AI reasoning in clear semantics. Conversely, recent work suggests LLMs can assist ontology development by digesting large corpora to suggest terms and relationships [15]. In this paper, we discuss how semantic interoperability via our ontology can support advanced applications, from AI-driven decision support (e.g. suggesting waste reduction strategies) to linked data platforms for sustainability reporting and certification. As an example, we incorporate the Upcycled Certified© standard (2025) into our model: Upcycled foods are an emerging valorization approach where surplus ingredients are turned into new consumer products. The *Upcycled Food Association's* standard provides criteria (e.g. at least 10% of ingredients from verified waste streams, documentation of amount diverted, etc.) that we translate into ontology terms (such as an *"Upcycled Food Product"* class linked via *derivedFrom* to a *Food Waste* source, and an *eligibleForCertification* property to flag products meeting the standard). By capturing such details, the ontology can facilitate traceability (tracking waste-derived ingredients), sustainability reporting (quantifying diverted waste), and semantic certification (digitally representing certifications and their criteria).

In summary, the Introduction has outlined the motivation and scope of our work. The following sections detail our methodology in constructing the ontology, present the resulting framework with examples of its classes and relations, discuss how it addresses prior challenges, and conclude with implications for research and practice.

2. Methods

We followed a structured methodology to design the Integrated Food Waste Ontology, emphasizing alignment with upper-level principles, reuse of existing ontologies, and inclusion of domain expert knowledge. The development process comprised several key steps:

Ontology Alignment and Upper-Level Design: We adopted Basic Formal Ontology (BFO**)** as the top-level framework to categorize entities and processes. All core classes in our ontology are positioned under appropriate BFO categories (e.g. *Food Waste* as a kind of Material Entity (continuant), *Food Waste Generation* as a Process (occurrent)). This BFO alignment ensures logical rigor and compatibility with many OBO ontologies. For example, instances of *Food Waste* can bear roles or qualities (such as a *waste role* or *edibility quality*), while *Food Waste Prevention* is modeled as a subtype of *planned process* (an intentional occurrent). Using BFO's continuant/occurrent distinction helped clearly separate physical waste materials from activities/events, a crucial design choice for representing dynamic waste flows.

Reusing Classes from OBO Foundry Ontologies: In line with best practices, we maximized reuse of existing ontology terms to avoid reinventing the wheel. FoodOn (OBO Food Ontology) was imported for general food product taxonomy and food material definitions. For instance, FoodOn's class *food material* was reused as a parent of our *Food Waste* class, indicating that food waste is still fundamentally food (just no longer intended for consumption). We also leveraged ENVO (Environment Ontology) for its definition of *waste material.* ENVO provides a class *"food waste"* defined as *"a waste material primarily composed of uneaten food removed from the food supply chain"*, which we adopted to ground our *Food Waste* class definition [8]. While this definition is suitable for redirecting *"food waste"* into non-edible processing streams, *"food waste"* itself becomes a misnomer when it is then upcycled into new food products, as the food is not then wasted. For instance, when spent grains from a brewing process are removed from the human food supply and sent to a hog farm, they are "wasted", however, when those same spent grains are "rescued" and upcycled into bread, then the item is no longer considered "*food waste"*.

Where available, related ENVO terms (e.g. *compost*, *anaerobic digestion process*) and RO (Relations Ontology) properties (e.g. *has part*, *derives from*) were utilized to maintain consistency with established vocabularies. Basic relations like *part-of* (for hierarchical supply chain relations) and *has output/has input* (for linking processes and material entities) follow RO's standard semantics.

Integrating Sustainability and Domain-Specific Ontologies: The ontology scope intersects with sustainability and climate domains, so we integrated terms from relevant ontologies:

- The **Sustainability Core Ontology (SCO)** [11] provided high-level concepts such as *Sustainability Objective* and *Indicator*, which we use to contextualize FLW reduction targets and metrics.
- **SDG Interface Ontology (SDGIO)** [12] terms were imported to link our classes to the UN Sustainable Development Goals. Notably, our *Food Waste Prevention* class is connected to *SDG Target 12.3 (food waste reduction)* as a related objective, enabling queries that relate ontology instances to SDG-aligned outcomes.
- **Sustainable Meat Supply Ontology (SuMSO)** [13] inspired how to handle supply chain stages and byproducts. For example, SuMSO's modeling of animal by-products guided our inclusion of *Inedible Portion* as a subclass of *Food Waste* (since bones, peels, etc. are "associated inedible parts" often counted in FLW [16].
- The **Sustainable Development and Climate Ontology (SDC)** [14] contributed terms to characterize environmental impacts (e.g. *Greenhouse Gas Emission*). We linked the outcome of *Food Waste Valorization* processes to potential impact mitigation (for instance, a *valorization process* of anaerobic digestion can be related to *biogas production* and GHG reduction, using SDC terms).

Incorporating Domain Definitions and Standards: We gathered authoritative definitions and modeling patterns from key domain resources to inform class definitions and relationships:

From FAO and USDA guidance, we ensured the ontology distinguishes "Food Loss" vs "Food Waste" contexts. We created subclasses or tags to indicate at what supply chain stage a waste event occurs (up to retail = food loss, retail/consumer = food waste) aligning with FAO's definitions. This allows users to classify an instance of *Food Waste* as *food loss* or *food waste* per the scenario, while maintaining a unified *Food Waste* concept for all material that exits the human food chain.

The **Food Loss & Waste Accounting and Reporting Standard (FLW Standard)** [16] shaped how we model quantities and destinations. The FLW Standard emphasizes quantifying the *mass of food and inedible parts removed from the supply chain* and tracking "where it goes" (destination). To represent this, we included an *Information Entity* class for *FLW Report* (subclass of IAO: dataset) which can use properties like hasQuantity (linked to a weight measurement data item) and hasDestination (linked to a controlled vocabulary of destinations such as *Donation*, *AnimalFeed*, *Landfill*, *Compost*). We incorporated the standard's specified destination categories (e.g., donated, co-product for animal feed, bio-based materials processing, composting/aerobic digestion, anaerobic digestion, land application, controlled combustion, sewer, landfill) as individuals or classes in a hierarchy of waste fates. This integration ensures that data annotated with our ontology can be directly used to produce FLW Standard-compliant reports and comparisons.

The **Upcycled Certified® Standard (2025)** [17] was analyzed to embed the concept of upcycled food products and their certification. From the standard, we extracted key criteria: (a) an upcycled product must contain a minimum percentage of ingredients derived from surplus or waste streams; (b) the source of those upcycled ingredients must be verified and traceable; (c) the producer should report the amount of waste diverted by the product annually, etc.. Accordingly, we modeled an Upcycled Food Product class (as a subtype of FoodOn's *food product*) with a restriction that some ingredient is derived from a Food Waste source (see Results for formal axioms). We also introduced an *eligibleForCertification* property that can link an *Upcycled Food Product* or *Ingredient* to a specific *Certification Scheme* individual (e.g., an instance representing the "Upcycled Certified" program). This allows the ontology to represent statements like "Product X *eligibleForCertification* UpcycledCertifiedProgram" when criteria are met—a stepping stone toward machine-readable certification and auditing frameworks.

Defining Novel Classes and Relationships: Where no suitable term existed in available ontologies, we created new classes and object properties, carefully documenting their definitions and intended usage. The top-level classes we introduced—*FoodWaste*, *FoodWasteGeneration*, *FoodWastePrevention*, *FoodWasteRecovery*, *FoodWasteValorization*, *UpcycledFoodProduct*—were defined with textual definitions drawn from literature and standards. For example, *FoodWaste* is defined as "food material originally produced for human consumption that is not ultimately consumed and is removed from the food supply chain", unifying FAO's notions of food loss/waste. We also coined domain-specific object properties to capture critical relationships:

- **hasWasteStream**: links a process to the *Food Waste* material stream it produces. This relation is used to assert, for instance, that a *food processing activity* or a *spoilage event* hasWasteStream some *FoodWaste* (the by-product or discarded food from that process).
- **derivedFrom**: a relation (aligned with RO's notion of derivation) that links a material or product to its source material. We use it to connect upcycled or recycled products to the original *Food Waste* input (e.g., "Bread beer is derivedFrom stale bread (a FoodWaste)"). This property aids in provenance tracking of valorized products.
- **eligibleForCertification**: a custom relation to indicate that an entity (e.g., a product or a process) qualifies for a certain certification or label. We apply this to model certification eligibility (for upcycled foods, organic labels, etc.), which typically depends on meeting criteria representable in the ontology (such as origin of ingredients, percentage thresholds, presence of monitoring processes, etc.).

  Each new term was created under a namespace for our ontology (provisionally "FWPO" – Food Waste Prevention Ontology) and, where possible, aligned as a sub-property or sub-class of an existing ontology term. For example, *hasWasteStream* can be considered a sub-property of a more general *hasOutput* property in RO, and *eligibleForCertification* relates to the concept of an entity having a *role* of being certified (we modeled certification as a *role* that a product can bear, following BFO's Role pattern).

Iterative Review and Expert Feedback: The ontology was developed iteratively, with continuous validation against the requirements gleaned from literature and stakeholder input. We cross-checked that known gaps identified by Lange et al. were addressed: e.g., ensuring *FoodWasteRecovery* and *FoodWastePrevention* classes exist (to cover rescue and reduction use cases), and that properties for *destination* and *cause* were included (to capture dynamic aspects) [2]. We also consulted informally with domain experts in food sustainability and ontology engineering to refine class scopes. For instance, stakeholder feedback emphasized the need to represent quantities and units (for reporting), so we linked our ontology to the OM (Ontology of Units of Measure) for standardized quantity values. The development adhered to OBO Foundry principles like openness, collaboration, and use of common relations, setting the stage for potential submission to the OBO library in the future.

By following these steps, we ensured the ontology is built on a strong foundation (BFO and existing ontologies) while richly encoding the domain-specific nuances of food waste generation and management. The next section (Results) describes the resulting ontology structure and gives examples of how the core classes and properties are implemented, including a snippet of the Turtle serialization of the ontology.

3. Results

3.1. Ontology Content and Structure

The Integrated Ontology for Food Waste Prevention, Recovery, and Valorization comprises a set of top-level classes and relations that together model the lifecycle of food waste and interventions on it. Figure 1 (see TTL excerpt below as a textual representation) outlines the major classes and their relationships. The ontology's central node is the class FoodWaste, representing the material entity of wasted food. *FoodWaste* is formally defined as a subclass of both foodon:FoodMaterial and envo:FoodWaste—indicating it is a food material by nature and a waste material by context. This aligns the class with the broader food ontology and environmental ontology definitions. All specific instances of discarded food (uneaten leftovers, spoiled produce, processing by-products, etc.) can be typed as *FoodWaste* in our framework. Four key process classes capture different stages and strategies related to food waste:

FoodWasteGeneration: a process in which food is rendered as waste. This class covers events or activities that cause food to exit the edible supply chain. It could be unintentional (e.g. spoilage, contamination) or intentional (e.g. discarding cosmetically imperfect produce). In the ontology, *FoodWasteGeneration* is a subclass of bfo:Process and is further characterized by an existential restriction: it hasWasteStream some FoodWaste. In other words, any instance of a *FoodWasteGeneration* process produces at least one output that is a FoodWaste material. This axiom provides a logical linkage between the process and the material outcome.

FoodWastePrevention: a process that aims to reduce or avoid the generation of food waste. This class is modeled as a subtype of bfo:Process, often a planned process, encompassing activities like inventory management, farm storage improvements, consumer education campaigns on food planning, etc. A *FoodWastePrevention* process does not necessarily produce a FoodWaste output (indeed, its goal is to minimize such output), so we do not constrain it with a hasWasteStream relation. Instead, we capture its intent via annotations or by linking it to a *Food Waste Reduction Goal* (an instance of *SustainabilityObjective*). For example, a campaign to improve household food storage can be an instance of *FoodWastePrevention* linked to a metric of waste reduction.

FoodWasteRecovery: a process of recovering or rescuing surplus food for human consumption or animal feed, rather than letting it become waste. This includes activities like food donation programs, gleaning leftover crops, redistributing unsold bakery items to food banks, etc. *FoodWasteRecovery* is an occurrent wherein *FoodWaste* (or *surplus food* that would otherwise become waste) is an input, and the output is re-purposed food that re-enters the supply chain (for humans or animals). We model this by linking *FoodWasteRecovery* to *FoodWaste* via either an input relation (e.g., RO:has_input) or a specialized property like *reclaims* (not formally defined in this paper). While we did not encode an explicit restriction in the snippet for simplicity, one could state: *FoodWasteRecovery subClassOf (hasInput some FoodWaste)*. This ensures any recovery process involves handling food that would otherwise be waste. The outcome of recovery might be classified as an *EdibleProduct* or *AnimalFeed* depending on use.

FoodWasteValorization: a process of converting food waste into new products or value, typically not for direct human food but for other purposes. This covers a range of valorization pathways: e.g. processing waste into animal feed ingredients, composting to create fertilizer, anaerobic digestion to produce energy (biogas), extracting chemicals or materials (like fibers, proteins) for industry, or even upcycling into new human food products (a special case overlapping with recovery). *FoodWasteValorization* is a subclass of bfo:Process. Similar to recovery, it would have *FoodWaste* as an input (the raw material being valorized) and yield one or more outputs, such as a *ValorizedProduct*. We introduced a general class *ValorizedProduct* (subclass of material entity) to encompass outputs like compost, biofuel, upcycled ingredients, etc., each of which can be further specified. The ontology thus can represent, for instance, an anaerobic digestion process instance that *has input* some food waste and *has output* some biogas (which would be an instance of a fuel class linked to energy ontologies). By including *FoodWasteValorization*, we cover what earlier ontologies did (many focused on valorization) but crucially, we place it in context with prevention and recovery classes, providing a fuller picture of FLW mitigation actions.

In addition to these main classes, we extended the ontology with supporting classes for destinations and causes to enable the dynamic, process-oriented descriptions noted as lacking in previous work. For destinations, we added a class hierarchy under *WasteDestination* (linked to the FLW Standard categories). Instances of FoodWaste can be associated with a destination via a data property or annotation (e.g., *hasDestination: Landfill* or *hasDestination: CompostingFacility*). For example, if a particular *FoodWasteGeneration* event is the dumping of produce, we can annotate that the *FoodWaste* generated has destination *Landfill*. For causal antecedents, we introduced an optional relation *hasCausalFactor*, which can point to a process or condition that precipitated the waste (modeled similarly to cause-effect in other ontologies). For instance, an instance of *FoodWasteGeneration* could *hasCausalFactor: PowerOutage* (an event leading to refrigerator failure and spoilage). This kind of representation addresses the noted blind spot of capturing *why* waste occurs, an important feature for analysis and prevention strategies.

3.2. TTL Implementation Excerpt

Below is an excerpt of the ontology in Turtle (TTL) syntax, illustrating the definition of core classes and properties discussed. This snippet (Listing 1) shows the creation of the *FoodWaste* class with links to FoodOn and ENVO, the definition of process classes with BFO alignment and some restrictions, the *UpcycledFoodProduct* class with a derivation constraint, and the object properties *hasWasteStream*, *derivedFrom*, and *eligibleForCertification*.

Listing 1: Excerpt from the ontology's Turtle representation, showing key classes and properties. Namespace prefixes for clarity: foodon : = Food Ontology, envo : = The Environment Ontology, bfo : = Basic Formal Ontology, : = our ontology's default namespace.

```
:FoodWaste a owl:Class ;
   rdfs:subClassOf foodon:FoodMaterial , envo:FoodWaste ;
   rdfs:label "Food Waste" ;
   rdfs:comment "Food material originally intended for human consumption removed from the food supply chain (includes 'food loss' and 'food waste')." .

:FoodWasteGeneration a owl:Class ;
   rdfs:subClassOf bfo:Process ;
   rdfs:label "Food Waste Generation" ;
   rdfs:subClassOf [ owl:onProperty :hasWasteStream ;
                owl:someValuesFrom :FoodWaste ] .

:FoodWastePrevention a owl:Class ;
   rdfs:subClassOf bfo:Process ;
   rdfs:label "Food Waste Prevention" .

:FoodWasteRecovery a owl:Class ;
   rdfs:subClassOf bfo:Process ;
   rdfs:label "Food Waste Recovery" .

:FoodWasteValorization a owl:Class ;
   rdfs:subClassOf bfo:Process ;
   rdfs:label "Food Waste Valorization" .

:UpcycledFoodProduct a owl:Class ;
   rdfs:subClassOf foodon:FoodProduct ,
                [ owl:onProperty :derivedFrom ;
                 owl:someValuesFrom :FoodWaste ] ;
   rdfs:label "Upcycled Food Product" .
:hasWasteStream a owl:ObjectProperty ;
   rdfs:domain bfo:Process ;
   rdfs:range :FoodWaste ;
   rdfs:label "has waste stream output" .
```

```
:derivedFrom a owl:ObjectProperty ;
  rdfs:label "derived from" .

:eligibleForCertification a owl:ObjectProperty ;
  rdfs:label "eligible for certification" .
```

In this snippet, we see, for example, that FoodWasteGeneration is defined as a process that by definition has some *FoodWaste* as output (via the owl:Restriction on *hasWasteStream*). Similarly, UpcycledFoodProduct is defined as a food product that is derived from some *FoodWaste*. These logical axioms allow a reasoner to infer relationships; e.g., if a particular product is linked via *derivedFrom* to a waste, the reasoner can classify that product as an *UpcycledFoodProduct*. This ontology uses if-then inferences to connect waste data to higher-level concepts. A proof-of-concept for this automated reasoning is provided (Listing 2). The property declarations show that *hasWasteStream* expects a process as domain and FoodWaste as range, while *derivedFrom* and *eligibleForCertification* are left more general (their domain/range can be any appropriate class—we kept them broad to allow use on different entity types, though in practice we may constrain them further in a future iteration).

Listing 2: Proof-of-concept example using existing ontology terms to support if-then inference reasoning.

Given: If a FoodWasteGeneration process
With a FoodWaste output
Whose hasDestination is UpcycledFoodProduct
Then one can infer: ?processInstance a: FoodWasteRecovery

3.3. Evaluation and Alignment with External Ontologies and Data

The ontology was evaluated for its coverage and alignment. All top-level classes have been mapped to at least one external reference or ontology ID to facilitate integration:

- **FoodWaste** is linked to ENVO:03600006 (Food waste) and contextualized with FAO's definition as a comment. This helps anyone using our ontology to see the correspondence with FAO's concept and ENVO's library.
- **FoodWastePrevention/Recovery/Valorization** are novel composite concepts, but they align with terms in sustainability discourse (e.g., prevention aligns broadly with waste minimization strategies in SDGIO, *recovery* with concepts in the EPA hierarchy, *valorization* with terms used in waste management ontologies). We included broad mappings where possible (for instance, an *exactMatch* link of *FoodWasteValorization* to a term in the EU Waste Framework Directive's categorization if available).
- **UpcycledFoodProduct** does not exist in other ontologies yet (to our knowledge), as it is emergent, but we anchor its definition in the Upcycled Food Association's standard. We anticipate linking this to forthcoming ontology efforts in the circular economy domain.

We also ensured that our object properties do not duplicate existing relations. *derivedFrom* is conceptually similar to RO's *derives_from* relationship, and we could make *:derivedFrom* a sub-property of that RO relation for clarity. The *eligibleForCertification* property is domain-specific; however, it relates to patterns of *has quality* or *has role* (the product has the role of being eligible for a certification). We considered modeling certifications as instances of *Role* that a product can acquire. In the current ontology, we opted for a direct property for simplicity, but in future work this could be refined using a role-based design (e.g., define an *UpcycledEligibleProductRole* in BFO's role hierarchy).

Finally, we tested the ontology with example individual instances to verify its logical consistency and usefulness. For instance, we created a demo individual :StaleBread of type *FoodWaste*, and :BeerFromBread of type *FoodProduct* with an assertion :derivedFrom :StaleBread. The reasoner (using OWL inference) successfully classified :BeerFromBread as an *UpcycledFoodProduct* due to our owl:Restriction axiom. Similarly, we instantiated :GleaningEvent1 as a *FoodWasteRecovery* and gave it an input some *FoodWaste*. We checked that all class constraints held and that no unintended inferences occurred (ensuring, for example, that a prevention process with no waste output isn't misclassified as a generation process). These evaluations proved that there is logical consistency, inference validation, and interoperability within the framework.

4. Discussion

Our proposed ontological framework addresses several of the key challenges and recommendations highlighted in prior research on food waste ontologies. In this section, we discuss how the ontology improves upon the state-of-the-art, the implications for interoperability and data integration, and the potential uses in AI and linked-data applications.

Comprehensive Coverage of FLW Interventions: A primary contribution of this work is filling the gap in representing food waste prevention and recovery alongside valorization. Previous ontologies overwhelmingly centered on downstream waste uses (valorization), neglecting upstream actions. By explicitly modeling *FoodWastePrevention* and *FoodWasteRecovery* as first-class concepts, our ontology realigns the knowledge representation with the food waste hierarchy endorsed by policy and sustainability experts. This means data about initiatives like food donation programs or waste-reducing innovations can be semantically annotated and linked to the same framework as data about composting or biogas production. For example, a city's food rescue operations (recovery) and its compost facilities (valorization) can both be represented and connected within our ontology, enabling a more holistic analysis of how the city is addressing FLW across the hierarchy of preferred options. This comprehensiveness is critical: reduction and recovery are often more impactful in reducing total waste, and having them in the ontology allows decision-support systems to reason about scenarios like "How much waste could be avoided by prevention vs. managed by recycling?" or "Which prevention measures target the causes of waste captured in our data?". Coverage of food and waste materials remains high while recovery pathways and consumer behaviors remain largely uncovered.

Process-Oriented, Dynamic Modeling: Our ontology places significant emphasis on processes and their attributes (time, location, cause, output), directly responding to the identified need for more dynamic descriptors in FLW ontologies. By representing waste generation events, linking waste flows to destinations, and capturing causal factors, we create opportunities for cross-sector traceability and analysis that were previously limited. Now, stakeholders can use the ontology to trace a waste stream from its origin (e.g., a farm or factory process) to its end fate (e.g., animal feed, landfill) in a standardized way. This level of detail supports systems-level interventions: for instance, policymakers could query a knowledge graph built on this ontology to find "all causal antecedents of retail food waste in region X" and discover patterns (like refrigeration failures or over-ordering) that suggest preventive actions. The ontology's capability to encode *w*aste as it happens (capturing flows and drivers) is a step towards the machine-readable representations that Lange et al. called for [2]. It enables what-if analyses and scenario modeling, e.g. if a certain causal factor is eliminated, how does that propagate to reduced waste quantities?

Improved Interoperability through Foundational Alignment: By anchoring our ontology in BFO and integrating classes from OBO Foundry ontologies (FoodOn, ENVO, etc.), we enhance term interoperability and encourage the ontology's use as a hub in a network of ontologies. This addresses the critique that earlier FLW ontologies lacked alignment with gold-standard frameworks and external ontologies. In practice, this means that data annotated with our ontology can interlink with databases using FoodOn (for food types), ENVO (for environmental contexts like storage conditions or waste treatment facilities), or other OBO ontologies. For example, consider a dataset of agricultural losses annotated with our ontology: *FoodWasteGeneration* events on farms can be linked to ENVO terms for environmental conditions (drought, pests), or to agroontologies for crop types, thereby enriching the dataset's context. Because we adhered to OBO principles of unique identifiers and open accessibility, our ontology can be readily extended or imported by others, fostering a community-driven evolution. We also plan to register the ontology in public repositories (like the OBO Foundry or AgroPortal) to maximize its visibility and reuse, aligning with recommendations that new FLW ontologies be developed within such foundry ecosystems.

Integration of Standards and Domain Knowledge: A notable feature of our framework is the incorporation of definitions and structure from domain standards (FLW accounting, upcycling certification). This ensures semantic consistency between data collection/reporting practices and the ontology model. Organizations that use the WRI FLW Standard for reporting, for instance, can map their terms (like "material type", "destination") directly to our ontology's classes and properties. This streamlines data conversion into a knowledge graph, preserving crucial distinctions like *edible vs. inedible portions*, or *donation vs. waste recycling*, which might otherwise be lost in a generic data model. Similarly, by encoding upcycling criteria, the ontology can become part of a semantic certification framework—enabling digital systems to automatically check if a product meets the ontological definition of *UpcycledFoodProduct* (e.g., has a waste-derived ingredient and relevant documentation) and perhaps flag it as eligible for certification. This kind of smart data infrastructure could reduce friction in sustainability certification processes and improve trust and transparency (consumers or auditors could trace the certified product's ingredients back to their waste sources via ontology links). Our work demonstrates a model for embedding *sustainability standards* into ontologies, which could be extended to other certifications (organic, fair trade, etc.), contributing to more standardized sustainability reporting.

Support for AI and LLM Applications: As AI, and specifically large language models (LLMs), become more prevalent in analyzing and generating insights from text and data, the presence of a formal ontology for food waste can significantly enhance their performance in this domain. LLMs on their own may have *knowledge* about food waste from training data, but that knowledge is unstructured and prone to misunderstanding ambiguous terms. By integrating our ontology into knowledge graphs or prompting frameworks, LLMs can be guided by a controlled vocabulary and relationships. For example, an LLM-based system could be connected to a knowledge graph of city-level waste data annotated with our ontology; the LLM, when asked a question like "How can our city reduce restaurant food waste?", can query the graph using the ontology structure (identifying which data points are instances of *FoodWasteGeneration* in restaurants, what causes are linked, etc.) to formulate a fact-based answer. This synergy is in line with suggestions that LLMs can assist in ontology curation, and we argue it works both ways: ontologies assist LLMs by providing semantic clarity and checkable facts. Our ontology could also facilitate machine reasoning combined with LLMs; for instance, a reasoner can infer that "upcycled beer" is an instance of *FoodWasteValorization*, and an LLM can then generate narratives or recommendations knowing the beer is part of waste mitigation strategies.

Additionally, the ontology's richness enables advanced analytics using AI/ML. Because it standardizes data, large datasets from different sources (farms, retailers, waste logs, etc.) can be merged and fed into machine learning models to predict waste patterns or optimize interventions. Feature engineering for such models becomes easier with a consistent ontology: features like "presence of a prevention process" or "cause type = cold chain failure" can be uniformly extracted across datasets. The ontology also lends itself to building simulation and optimization tools (e.g., agent-based models of food supply chains as seen in Baena *et al.* 2020 could use our ontology to ensure agents share a common understanding of resources and waste flows) [18].

Limitations and Future Work: While our ontology provides a broad and integrated view, there are areas for further development. One limitation is the current level of granularity: many detailed subclassifications (for example, specific *FoodWastePrevention* techniques or finer distinctions of *FoodWaste* by type) are not exhaustively enumerated in the ontology. We intentionally focused on the upper level; however, domain experts could extend the ontology downward (e.g., adding subclasses for *HouseholdFoodWaste* vs *SupplyChainFoodWaste*, or specific valorization processes like *AerobicComposting* vs *AnaerobicDigestion*). Our use of general relations like *eligibleForCertification* could be refined by converting certain properties to *roles* or linking to existing certifications ontologies (if they emerge). We also acknowledge that the ontology's utility will depend on adoption by various stakeholders, which in turn may depend on demonstrating its value in pilot projects or tools. Ensuring the ontology stays updated with evolving terminology (for instance, new upcycling methods or emerging policies) will require a governance process, potentially under an open community or a standards body.

Another challenge is data integration effort: existing datasets need mapping to the ontology. This can be labor-intensive, but tools like mapping dictionaries or automated NLP (possibly assisted by LLMs) could help tag legacy data with ontology IDs. Our framework is at the conceptual level; in the future, we plan to develop user-friendly extensions or templates for different user groups (e.g., a simplified model for food businesses to track waste, or an extension focusing on household waste behaviors).

Finally, as noted in the literature, building such ontologies should be a collaborative effort. Our work sets the stage, and we invite contributions from others—be it aligning this with the next version of FoodOn (which is continually expanding) or integrating with initiatives like the proposed ontology by Weber et al. (2023) that was mentioned as a comprehensive approach [19]. Through collaboration, the ontology can evolve to truly become the "unifying data language" for food waste that multiple authors have identified as crucial for accelerating progress on this global challenge.

5. Conclusion

In this paper, we presented a full-spectrum ontological framework for food waste prevention, recovery, and valorization, designed in alignment with BFO and leveraging the rich ecosystem of OBO Foundry and sustainability ontologies. The ontology fills critical gaps by formally representing not just food waste materials, but also the key processes and interventions across the food waste mitigation hierarchy. By integrating terms and patterns from FoodOn, ENVO, SCO, SDGIO, SuMSO, SDC, and by incorporating definitions from FAO, USDA, the FLW Standard, and the Upcycled Food Standard, our framework achieves a high degree of semantic alignment with real-world knowledge and practices.

The resulting ontology enables stakeholders to connect previously siloed data on food waste – from on-farm losses to retail waste bins to post-consumer waste management – under a unified schema. This interoperability is a foundational step toward improved data sharing and analytics for FLW, as emphasized by recent research. Moreover, our inclusion of dynamic process attributes (cause, time, destination) and strategic interventions (prevention, recovery actions) means that the ontology can support more nuanced analyses and solutions development than static taxonomies of waste. Users can represent complex scenarios (e.g., "if X tons of food waste at wholesale are diverted to food banks, how does that impact overall waste?") in a queryable, machine-understandable way.

We demonstrated how the ontology can capture novel concepts like upcycled foods and certification criteria, pointing towards a future where sustainability certifications and reporting are facilitated by semantic web technologies. Tools built on this ontology could automate the tracking of how much waste was avoided or repurposed, providing evidence for sustainability claims. Furthermore, as artificial intelligence and large language models become integrated into decision-making, having a robust ontology ensures that these systems operate on a solid knowledge foundation, reducing ambiguity and increasing trust in their outputs.

Our work contributes a modular, extensible ontology that ontology engineers, sustainability informatics researchers, and food system scientists can build upon. We envision this ontology being used to create knowledge graphs for smart food waste management systems, to harmonize data for international FLW monitoring efforts, and to serve as a basis for educational tools that teach about food waste flows and solutions. We also see potential in coupling this ontology with others (e.g., supply chain ontologies, consumer behavior ontologies) to address the food waste issue from production to consumption with a systems perspective.

In conclusion, the Integrated Ontological Framework for Food Waste Prevention, Recovery, and Valorization offers a timely and comprehensive semantic infrastructure aligned with best practices and domain needs. It addresses previously identified shortcomings by covering the full range of FLW activities, embedding process dynamics, and ensuring interoperability and reusability. We hope this work will spark further collaboration and ontology development in the food sustainability domain. By enabling data to "speak the same language" across disciplines and sectors, such ontologies can help unlock new insights and drive informed actions towards the shared goal of reducing food waste and building more sustainable food systems.

Acknowledgements

The authors would like to thank Brian Roe and Karthik Jain for their conceptual contributions to earlier versions of this work.

Declaration on Generative AI

We used Brave Leo, Google Gemini-2.5, and ChatGPT-5.2 for conceptual clarification, content synthesis, and language refinement. All content was reviewed by the authors, who take full responsibility.

References

[1] US Environmental Protection Agency. Sustainable Management of Food Basics. US EPA. URL: https://www.epa.gov/sustainable-management-food/sustainable-management-food-basics (2015).

[2] Lange, M. C., Li, R. & Apolzan, J. W. Ontologies relevant for improving data interoperability for food loss and waste: A review and research agenda. Cleaner and Responsible Consumption. 19, (2025).

[3] Borghini, A. & Piras, N. Semantic and philosophical approaches for advancing the identification and measurement of food waste. Nat Food 6, 547–552 (2025).

[4] Food and Agriculture Organization of the United Nations. Food Loss and Waste and Value Chains: Learning Guide. (Food & Agriculture Org., 2019).

[5] Food and Agriculture Organization of the United Nations. Voluntary Code of Conduct for Food Loss and Waste Reduction. (Food & Agriculture Org., 2022).

[6] Basic Formal Ontology (BFO). Available: https://basic-formal-ontology.org/

[7] Dooley DM, Griffiths EJ, Gosal GS, Buttigieg PL, Hoehndorf R, Lange MC, et al. FoodOn: a harmonized food ontology to increase global food traceability, quality control and data integration. NPJ Sci Food. 2018;2: 23.

[8] Buttigieg, P. L. et al. The environment ontology: contextualising biological and biomedical entities. J. Biomed. Semantics 4, 43 (2013).

[9] Relations Ontology. [cited 30 Jan 2026]. Available: https://bioportal.bioontology.org/ontologies/OBOREL

[10] Information Artifact Ontology. [cited 30 Jan 2026]. Available: https://bioportal.bioontology.org/ontologies/IAO

[11] Sustainability Core Ontology - Segment U. [cited 21 Dec 2025]. Available: https://bioportal.bioontology.org/ontologies/SCO-U

[12] SDG-InterfaceOntology/sdgio: The repository for the Sustainable Development Goals Interface Ontology. In: GitHub [Internet]. [cited 19 Dec 2025]. Available: https://github.com/SDG-InterfaceOntology/sdgio

[13] SuMSO. [cited 21 Dec 2025]. Available: https://biodivportal.gfbio.org/ontologies/SUMSO

[14] Sustainable-Development-and-Climate-SDC-ontology: The Sustainable Development and Climate (SDC) ontology represents the types and interrelations of entities that are outlined in the targets and indicators of the United Nations Sustainable Development Goals (U.N. SDGs). Github; Available: https://github.com/hbabaie1/Sustainable-Development-and-Climate-SDC-ontology

[15] Fathallah, N., Staab, S., Alergawy, A. LLMs4Life: Large Language Models for Ontology Learning in Life Sciences. [cited 15 Dec 2025]. Available: https://arxiv.org/html/2412.02035v1

[16] Hanson, C. et al. Food Loss and Waste Accounting and Reporting Standard. (2016).
[17] Upcycled Food Association. Defining Upcycled Foods A Definition for Use Across Industry, Government, and Academia. Harvard Center for Health Law and Policy Innovation. URL: https://chlpi.org/wp-content/uploads/2013/12/Upcycled-Food_Definition.pdf.
[18] Baena, B. et al. Adapting food supply chains in Smart Cities to address the impacts of COVID19 a case study from Guadalajara metropolitan area. in 2020 IEEE International Smart Cities Conference (ISC2) 1–8 (IEEE, 2020). doi:10.1109/isc251055.2020.9239076.
[19] Weber, M., Buche, P., Ibanescu, L. et al. PO2/TransformON, an ontology for data integration on food, feed, bioproducts and biowaste engineering. npj Sci Food 7, 47 (2023). https://doi.org/10.1038/s41538-023-00221-2