ID: 738

A knowledge graph-based agricultural information architecture for sustainable decision support

Cevher Özden

Çukurova University, Faculty of Arts and Sciences, Department of Computer Sciences, 01250, Adana, Türkiye.

Abstract

Sustainable agricultural decision support increasingly depends on the ability to connect different kinds of knowledge rather than simply accumulate more data. Crop characteristics, soil properties, climate conditions, water requirements, pest and disease information, farm practices, regional constraints, and sustainability indicators all enter into agricultural decisions. In many existing digital agriculture systems, however, these resources are stored in separate structures and described with different terms and assumptions. This fragmentation reduces interoperability and reuse, and it also makes recommendations harder to explain. Prediction and monitoring tools are useful, but they often leave the relationships among crops, soils, environmental conditions, management practices, risks, and decision outcomes only weakly represented.

This study develops a knowledge graph-based agricultural information architecture for sustainable decision support. The architecture represents agricultural decision knowledge through connected entities such as crop, soil, climate condition, water resource, disease, pest, farm practice, region, sustainability indicator, and recommendation. These entities are linked through relations such as requires, is suitable for, increases risk of, reduces, is affected by, and is recommended for. Representing agricultural knowledge as a graph makes it possible to support structured queries, semantic integration, evidence tracing, and interpretable recommendations.

The contribution of the study is mainly architectural. It moves the discussion from isolated repositories and model-centered prediction tools toward relational knowledge infrastructures. Such an infrastructure can support crop-soil suitability analysis, irrigation advisory, disease-risk interpretation, sustainable practice recommendation, and region-specific agricultural planning. From a Management Information Systems perspective, the study treats knowledge graphs as a reusable semantic backbone for agricultural information systems. From a Computer Engineering perspective, it provides a graph-based structure for integrating heterogeneous agricultural data and supporting explainable, query-driven, and sustainability-oriented decision intelligence.

Keywords: *Knowledge graph, Agricultural information architecture, Sustainable agriculture, Decision support systems, Semantic integration, Relational decision intelligence, Smart farming*

Introduction

Digital agriculture now draws on a wide ecosystem of data sources: sensors, satellite observations, weather services, farm records, soil databases, crop calendars, disease reports, mobile advisory tools, and institutional datasets. The problem is not that these resources are unavailable. In many cases, the relevant information already exists. The difficulty is that it is scattered across platforms and described through different formats, labels, and assumptions. As a result, agricultural decision-support systems often have access to data without having a sufficiently connected knowledge structure for reasoning, explanation, and reuse.

This issue becomes especially visible in sustainability-oriented agriculture. A decision about irrigation, crop choice, disease prevention, or input use is rarely driven by a single variable. It usually depends on a combination of crop requirements, soil characteristics, local climate, water availability, pest and disease pressure, management practices, regional constraints, and environmental indicators. Earlier smart farming research has already shown that big data applications influence not only field-level production but also the wider agri-food chain. Yet the value of these data depends on how they are organized, shared, and converted into decision support (Wolfert et al., 2017).

Knowledge graphs offer a useful response to this design problem. They represent entities and relationships in a structured, machine-readable, and semantically meaningful form. Hogan et al. describe knowledge graphs as graph-based structures that support the use of diverse, dynamic, and large-scale collections of data. Ji et al. similarly emphasize their role in knowledge representation, acquisition, completion, and knowledge-aware applications (Hogan et al., 2021; Ji et al., 2022).

This paper does not introduce another agricultural prediction model. Its starting point is more basic: before prediction outputs can become reliable decision support, agricultural knowledge must be organized in a form that makes relationships explicit. For this purpose, the study proposes a knowledge graph-based agricultural information architecture. The architecture connects core entities such as crop, soil, climate condition, water resource, disease, pest, farm practice, region, sustainability indicator, and recommendation. Once these entities and their relations are represented explicitly, the system can support structured querying, semantic integration, evidence tracing, and interpretable decision support.

The study contributes in three ways. First, it treats agricultural decision support as a problem of relational knowledge organization, not merely as a data analytics task. Second, it proposes a domain-oriented architecture that links heterogeneous agricultural knowledge sources through a knowledge graph. Third, it illustrates how this architecture can support practical decision-support scenarios, including crop-soil suitability analysis, irrigation advisory, disease-risk interpretation, and sustainable practice recommendation.

Problem Background

Agricultural decision-support systems work with heterogeneous data. These data may include numerical sensor measurements, categorical soil descriptions, remote sensing products, meteorological observations, expert rules, farm management records, pest and disease descriptions, and policy-related sustainability indicators. Heterogeneity itself is manageable. The more difficult issue is semantic disconnection. The same concept may be named differently across systems, while concepts that are logically related may remain hidden in separate databases.

This semantic fragmentation weakens interoperability. A crop database may define water requirements; a soil database may describe drainage capacity, pH, and texture; a weather service may report rainfall anomalies; and an advisory tool may issue an irrigation recommendation. If these resources are not connected through shared concepts and explicit relations, the system cannot easily explain why a recommendation fits a particular crop, soil type, region, and climate condition. This concern is consistent with the digital agriculture literature, which argues that technical innovation needs to be examined together with data infrastructures, institutions, knowledge systems, and socio-technical arrangements (Klerkx et al., 2019).

A second problem is low reusability. Many agricultural information systems are designed for a narrow task such as yield prediction, disease identification, irrigation scheduling, or farm monitoring. They may work for their immediate purpose, but their knowledge structures are rarely reusable across other decision-support tasks. Knowledge used for crop-soil suitability, for example, may also be relevant for irrigation planning, disease-risk assessment, and sustainable input management. Without a relational structure, this knowledge remains trapped inside task-specific applications.

A third limitation is explainability. Data-driven platforms can generate recommendations, but they often say little about the evidence behind them. In agricultural decision-making, this is not a minor problem. Farmers, advisors, cooperatives, and public institutions may need to know which crop requirement was considered, which soil property mattered, which climate condition increased the risk, or which sustainability objective was affected. Ontology-based agricultural decision-support systems, such as AgrODSS, show the value of semantic models for representing plant disease and pest knowledge in a machine-understandable form and for providing evidence-based explanations (Alharbi et al., 2024).

A fourth issue concerns the absence of a reusable semantic backbone in many agricultural platforms. Agricultural knowledge organization systems already exist. AGROVOC, for instance, is a multilingual controlled vocabulary and linked data resource for food and agriculture, developed to improve access and interoperability across domains and languages (Subirats-Coll et al., 2022). Still, many decision-support platforms do not use such semantic resources as part of their internal information architecture.

These problems point to the same conclusion: sustainable agricultural decision support requires more than data collection and model development. It needs an information architecture that can connect heterogeneous knowledge sources, represent relations explicitly, and make recommendations interpretable.

Material and Methods

The architecture proposed in this study is built around an Agricultural Knowledge Graph. The graph acts as a semantic layer between raw data sources and decision-support applications. Instead of keeping agricultural information in isolated tables, files, or documents, the architecture represents domain knowledge as connected entities and relations.

At the core of the architecture are several entity types that commonly appear in agricultural decision-making:
Crop: wheat, tomato, maize, chickpea, potato, citrus.
Soil: clay loam, sandy soil, salinity condition, pH level, organic matter.
Climate Condition: rainfall anomaly, high humidity, heat stress, frost risk.
Water Resource: irrigation source, water availability, water requirement.
Pest and Disease: fungal disease, insect pest, disease-risk condition.
Farm Practice: drip irrigation, crop rotation, fertilization, pesticide application.
Region: province, district, agroecological zone, farm location.
Sustainability Indicator: water-use efficiency, soil health, input intensity, environmental risk.

Recommendation: irrigation advice, crop suitability advice, disease-risk warning, sustainable practice suggestion.

These entities gain decision-support value through their relations. Examples include:

Crop requires Water Requirement.

Crop is suitable for Soil Type.

Climate Condition increases risk of Disease.

Farm Practice reduces Water Use.

Region has Soil Property.

Recommendation is based on Evidence.

Sustainability Indicator is affected by Farm Practice.

This relational structure is the main difference between a conventional database and a knowledge graph. A database can store crop, soil, and weather records. A knowledge graph can additionally represent why these records are connected and how their relationships support a recommendation. Sustainable agricultural decision support needs this second capability, because it requires contextual interpretation as well as data retrieval.

The architecture can be organized into four layers.

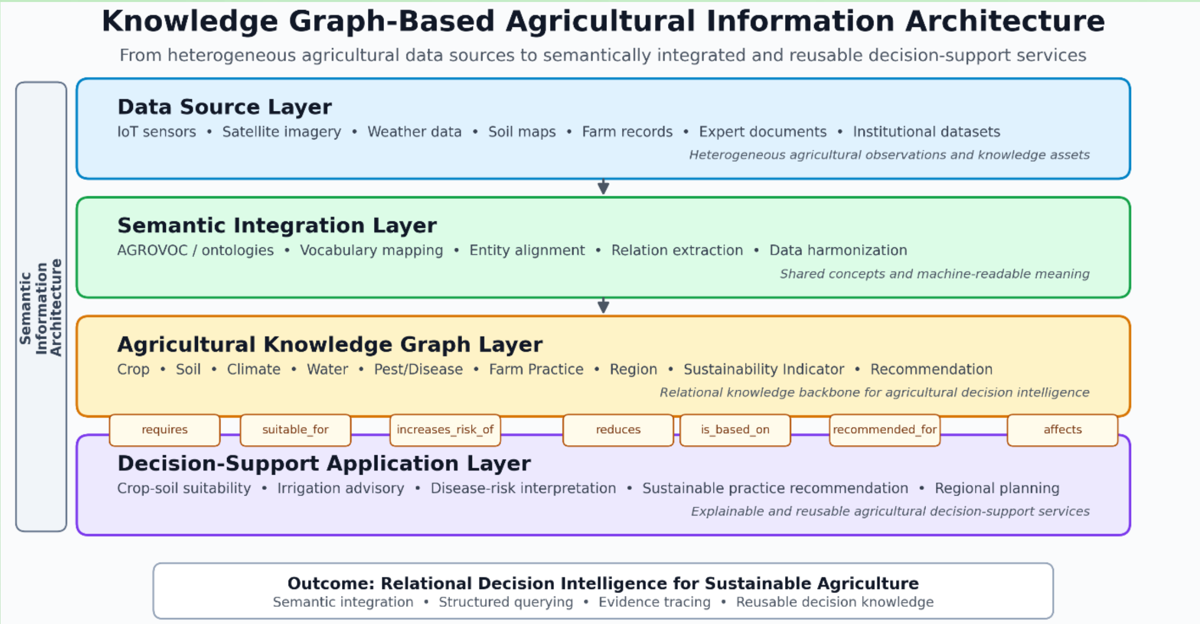

Figure 1. Knowledge Graph-Based Agricultural Information Architecture for Sustainable Decision Support

Data Source Layer: This layer contains raw and semi-structured agricultural data, including sensor measurements, satellite-derived indicators, weather records, soil maps, crop databases, expert documents, farm records, and institutional datasets.

Semantic Integration Layer: This layer maps heterogeneous data sources onto shared agricultural concepts. Controlled vocabularies and linked data resources, such as AGROVOC, can support semantic consistency and interoperability across systems (Subirats-Coll et al., 2022).

Agricultural Knowledge Graph Layer: This is the central layer of the architecture. It represents agricultural entities and their relations as a graph, connecting crops, soils, climate conditions, water resources, diseases, farm practices, regions, sustainability indicators, and recommendations.

Decision-Support Application Layer: This layer uses the graph to support advisory tasks such as crop-soil suitability analysis, irrigation planning, disease-risk interpretation, sustainable practice recommendation, and regional agricultural planning.

Table 1. Simplified entity-relation structure

Entity	Example	Relation	Decision-Support Value
Crop	Tomato	requires -> Water need	Supports irrigation planning
Soil	Clay loam	suitable_for -> Crop	Supports suitability analysis
Climate Condition	High humidity	increases_risk -> Disease	Explains disease risk
Farm Practice	Drip irrigation	reduces -> Water use	Supports sustainability advice
Region	Adana	has_climate -> Profile	Supports local advisory
Recommendation	Irrigation warning	based_on -> Evidence	Enables explainable support

The knowledge graph is therefore not only a storage layer. It is a reasoning-ready information structure. It allows a decision-support system to ask questions such as: Which crops match this soil and climate profile? Which disease risks increase under recent humidity and temperature conditions? Which practices reduce water use without raising disease risk? Which recommendations are supported by which evidence?

Results

The architecture can support several agricultural decision-support scenarios. Figure 2 presents an illustrative agricultural knowledge graph that connects crop, soil, climate, water, disease, practice, region, sustainability indicator, and recommendation entities through semantically meaningful relations.

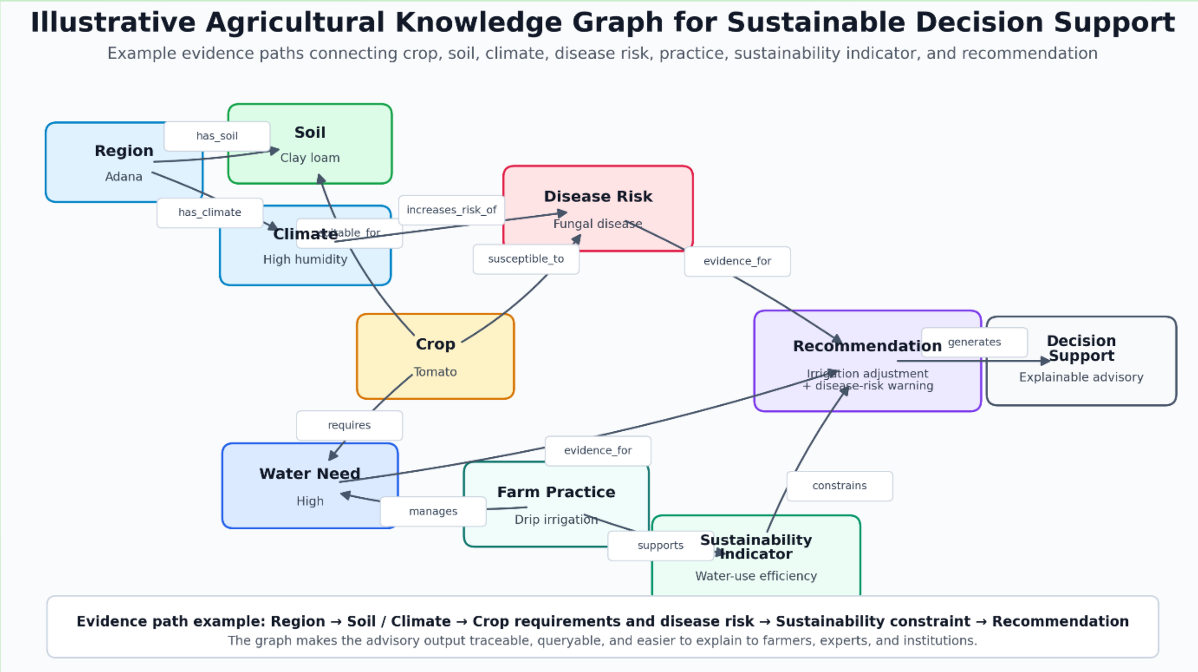

Figure 2. Illustrative Agricultural Knowledge Graph for Sustainable Decision Support

The following scenarios show how graph-based knowledge representation can improve interpretation and reuse.

Crop-Soil Suitability

A farmer or advisor may ask whether tomato is suitable for a specific field. A conventional system might consult a crop table and a soil table separately. A knowledge graph-based system can connect crop requirements, soil texture, pH level, salinity condition, drainage status, regional climate, and previous production constraints in one relational structure.

A graph-based query may follow a logic such as: find crops that are suitable for clay loam soil, tolerate the local pH range, match the regional temperature profile, and do not exceed available irrigation capacity.

The important point is that the answer is not just a label such as suitable or unsuitable. The system can show the supporting relations: crop-soil compatibility, water requirement, climate match, and regional constraints. This improves transparency for agricultural advisors and makes the recommendation easier to justify.

Irrigation Advisory

Irrigation decisions are shaped by several interacting factors, including crop growth stage, crop-specific water requirement, soil moisture, rainfall, evapotranspiration, irrigation method, and available water resources. In many agricultural information systems, these variables are stored or evaluated separately. A knowledge graph can bring them together within a reusable relational structure.

For example, tomato has a high water requirement during fruit development. If the region has received limited rainfall in recent days, the soil has only moderate water retention capacity, and the farm uses drip irrigation, the system may recommend irrigation. However, the recommendation should not stop at saying that irrigation is needed. It should also indicate that the amount of water should be adjusted carefully to avoid unnecessary use.

This scenario shows why sustainability indicators matter. The recommendation is not limited to a binary output such as irrigate or do not irrigate. It can also be linked to water-use efficiency, resource conservation, and environmental risk. In this sense, the graph supports both production-oriented and sustainability-oriented decision-making.

Disease-Risk Interpretation

Disease-risk warnings are often difficult to explain because they depend on interactions among humidity, temperature, rainfall, crop susceptibility, pest and disease knowledge, and farm management practices. Ontology-based systems in agriculture have shown that semantic models can support plant disease and pest identification by representing domain knowledge in machine-understandable form (Alharbi et al., 2024).

In the proposed architecture, a disease-risk warning can be generated and explained through graph relations. High humidity may increase the risk of fungal disease. Tomato may be susceptible to that disease under humid conditions. Recent rainfall may further support disease development, and the field may also have a previous disease record. When these pieces of evidence are connected, the system can explain why the disease risk is considered high.

This is more useful than a black-box alert. Instead of only presenting a warning, the system can show the evidence path behind it. Such evidence paths are particularly valuable for agricultural extension services, where recommendations need to be understandable, traceable, and defensible.

Sustainable Practice Recommendation

Sustainability-oriented agricultural decisions often involve trade-offs. A practice may reduce water use but increase labor requirements. Another may lower disease pressure but require additional management effort or chemical input. A knowledge graph can represent these trade-offs explicitly by connecting practices, expected benefits, possible risks, and sustainability indicators.

For example, drip irrigation may reduce water consumption, crop rotation may reduce disease pressure, and excessive nitrogen application may increase environmental risk. A decision-support system should therefore recommend a practice package that balances water efficiency, disease prevention, input intensity, and environmental protection.

The strength of the graph-based approach is that it can connect each practice to multiple sustainability indicators. This makes the architecture suitable for systems that need to move beyond productivity alone and include environmental, resource-efficiency, and long-term sustainability concerns.

Discussion

The MIS contribution of this study lies in a simple but often overlooked point: agricultural decision support is also an information architecture problem. In many digital agriculture projects, most attention goes to data collection, dashboards, and predictive models. These components are useful, but they do not by themselves create a reusable decision-support infrastructure. A sustainable system also needs a semantic structure through which data, expert knowledge, and decision logic can be connected and interpreted together.

From a Computer Engineering perspective, the knowledge graph provides a graph-based structure for integrating heterogeneous agricultural data. This matters because agricultural knowledge is not only numerical. It includes categories, textual descriptions, expert rules, regional classifications, spatial entities, temporal observations, and domain-specific concepts. Knowledge graphs are well suited to this complexity because they can represent entities, relations, schemas, context, and queryable knowledge structures within the same framework (Hogan et al., 2021; Ji et al., 2022).

From the perspective of relational decision intelligence, the architecture shifts attention from isolated prediction to connected reasoning. A machine-learning model may estimate disease risk or water demand, but a knowledge graph can place that output within a broader decision context. It can link the prediction to crop susceptibility, local climate, soil condition, management practice, sustainability target, and advisory recommendation. This makes the decision-support process more interpretable, reusable, and easier to justify.

The architecture also has practical value for agricultural institutions. It can support advisory platforms where knowledge from different departments, data providers, and expert groups needs to be integrated. It can also reduce the repeated development of disconnected systems for separate agricultural problems. Instead of creating isolated knowledge structures for irrigation, disease management, crop selection, and sustainability planning, institutions can build a shared agricultural knowledge graph and reuse it across services.

The study also has limitations. It presents a conceptual architecture and does not implement a complete agricultural knowledge graph. Future studies should operationalize the architecture using real datasets, domain ontologies, graph databases, and query examples. Further work may also evaluate the architecture with agricultural experts, compare it with conventional relational database designs, and test its usefulness in domains such as irrigation advisory, plant disease management, crop suitability analysis, and regional sustainability planning.

Conclusion

This study proposed a knowledge graph-based agricultural information architecture for sustainable decision support. Its main argument is that agricultural decision-support systems should not be designed only as data repositories, dashboards, or prediction tools. They should also function as relational knowledge infrastructures that connect crops, soils, climate conditions, water resources, pests, diseases, farm practices, regions, sustainability indicators, and recommendations within a shared semantic structure.

The architecture places the Agricultural Knowledge Graph at the center of the decision-support process. By representing agricultural entities and their relationships explicitly, it supports semantic integration, structured querying, evidence tracing, and interpretable recommendations. In practical terms, it can be used for crop-soil suitability analysis, irrigation advisory, disease-risk interpretation, sustainable practice recommendation, and region-specific planning.

For Management Information Systems, the study positions knowledge graphs as a semantic backbone for agricultural information systems. For Computer Engineering, it offers a graph-based structure for integrating heterogeneous agricultural data and supporting query-driven decision intelligence. For sustainable agriculture, it provides a way to connect production decisions with environmental and resource-efficiency indicators.

The main message is straightforward: digital agriculture should move beyond the ambition to collect more data. The next step is to organize agricultural knowledge relationally, so that decision-support systems can reason across crops, environments, practices, and sustainability goals.

Declarations

Ethical Approval Certificate

Not applicable. This study is a conceptual framework study and did not involve human participants, animals, clinical procedures, experimental interventions, or survey-based data collection requiring ethical approval.

Author Contribution Statement

Cevher Özden: Conceptualization, methodology, literature review, framework design, visualization, writing the original draft, review and editing.

Fund Statement

This work received no specific grant from any funding agency, commercial organization, or non-profit institution.

Conflict of Interest

The author declares no conflict of interest.

Acknowledgments

The author did not receive any administrative, technical, or material support requiring acknowledgment.

References

Alharbi, A. F., Alosaimi, W., Alharbi, A., & Alyoubi, K. H. (2024). An ontology-based agriculture decision-support system with an evidence-based explanation model. Smart Agricultural Technology, 9, Article 100568. https://doi.org/10.1016/j.atech.2024.100568

Hogan, A., Blomqvist, E., Cochez, M., D'amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Labra Gayo, J. E., Navigli, R., Neumaier, S., Ngonga Ngomo, A.-C., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2021). Knowledge graphs. ACM Computing Surveys, 54(4), Article 71. https://doi.org/10.1145/3447772

Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. IEEE Transactions on Neural Networks and Learning Systems, 33(2), 494–514. https://doi.org/10.1109/TNNLS.2021.3070843

Klerkx, L., Jakku, E., & Labarthe, P. (2019). A review of social science on digital agriculture, smart farming and Agriculture 4.0: New contributions and a future research agenda. NJAS: Wageningen Journal of Life Sciences, 90–91, Article 100315. https://doi.org/10.1016/j.njas.2019.100315

Subirats-Coll, I., Zeng, M. L., Li, D., & Nogales, E. G. (2022). AGROVOC: The linked data concept hub for food and agriculture. Computers and Electronics in Agriculture, 196, Article 105965. https://doi.org/10.1016/j.compag.2020.105965

Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M.-J. (2017). Big data in smart farming—A review. Agricultural Systems, 153, 69–80. https://doi.org/10.1016/j.agsy.2017.01.023