

Semantic Knowledge Graphs, G³, and Sustainable AI: Aligning Innovations with ESG Objectives

By Dinis Cruz and ChatGPT Deep Research, 2025/06/24

Introduction

In an era when organizations are increasingly accountable for the **environmental, social, and governance (ESG)** impacts of their technology, innovative approaches in data and AI engineering are emerging as key enablers of sustainable and ethical practices. This white paper, co-authored by Dinis Cruz and ChatGPT Deep Research, explores how **semantic knowledge graphs** and the concept of **G³ (Graphs of Graphs of Graphs)** pioneered in Cruz's research align with current efforts to make IT and AI systems more sustainable and ESG-compliant. We examine how these practices intersect with themes like **carbon-efficient computing, ethical AI, stakeholder transparency, open-source collaboration, responsible data governance, and AI explainability**. In doing so, we draw connections between Cruz's work and broader sustainable AI initiatives, providing conceptual mappings and actionable takeaways for practitioners and stakeholders.

Dinis Cruz's Semantic Knowledge Graphs and G³ Approach

Dinis Cruz's work centers on harnessing semantic knowledge graphs to manage and interconnect complex domains of information. **Semantic knowledge graphs** are structured representations of knowledge that link entities through well-defined relationships, adding context and meaning to data. For example, a knowledge graph might encode facts like *"Paris is the capital of France"*, connecting the entities *Paris* and *France* via a meaningful relationship ¹. By contextualizing data in this way, knowledge graphs enable more accurate reasoning and retrieval for AI systems ². Cruz's engineering practice pushes this concept further through **G³, or "Graphs of Graphs of Graphs,"** which advocates connecting multiple graphs and ontologies rather than forcing a single dominant hierarchy. As Cruz notes, G³ means having "ontologies of ontologies and taxonomies of taxonomies" – there is no one master uber-ontology, but instead a way to **communicate and connect multiple graphs, domains and ultimately different points of view** ³. In essence, G³ is about federating knowledge sources and perspectives, allowing systems to remain flexible, adaptive, and inclusive of diverse viewpoints.

This approach addresses a key challenge in knowledge management: different teams or sectors often develop their own taxonomies or data schemas. Rather than impose a single rigid structure, G³ enables interoperability among these structures, reflecting Cruz's vision that meaning-making in AI should be a **"shared civic responsibility"** supported by transparency and critique (as he has discussed in community dialogues). By supporting multiple coexisting ontologies held together through transparent linking, Cruz's graph-based approach inherently values **pluralism, adaptability, and trust** in knowledge systems. These values strongly resonate with ESG principles, as we will explore in the following sections.

Before diving into each ESG dimension, it's useful to frame what ESG means in the context of sustainable AI and IT systems. ESG encompasses three facets: **Environmental** (e.g. carbon footprint and resource use), **Social** (e.g. ethical practices, inclusivity, and societal impact), and **Governance** (e.g. transparency, accountability, data governance). Recent thought leadership emphasizes that “sustainable AI” involves aligning AI development with responsible, ethical principles to promote positive outcomes across all three ESG pillars ⁴. In practice, this means minimizing the environmental impact of computing (for example, using energy-efficient models and hardware) ⁵, ensuring AI systems are fair, inclusive, and respect human rights ⁶, and instituting strong governance around transparency, privacy, and accountability in AI ⁷. These goals are increasingly echoed by policymakers and industry initiatives. The European Union's draft **AI Act**, for instance, makes it a priority that AI systems deployed in the EU are “safe, transparent, traceable, non-discriminatory and environmentally friendly” ⁸. Likewise, organizations such as the **Green Software Foundation** argue that **responsible AI must consider carbon emissions alongside social ramifications**, advocating for lifecycle accountability and standards to measure and reduce AI's environmental footprint ⁹. With this context in mind, we analyze how Cruz's semantic graph approach contributes to each of these areas.

Carbon-Efficient Computing and Environmental Sustainability

One of the most pressing ESG concerns in IT today is the **carbon footprint** of computing, especially AI model training and inference. The energy consumption of large-scale AI has skyrocketed, leading to alarmingly high emissions. Researchers found that training a single medium-sized AI model (using neural architecture search) could generate emissions roughly equivalent to **626,000 pounds of CO₂** – on the order of the lifetime emissions of five average American cars ¹⁰. This has spurred calls for “Green AI” that prioritizes energy efficiency and carbon reduction over raw performance gains. Cruz's work on knowledge graphs, while not explicitly an environmental project, aligns with the need for more **carbon-efficient computing** in several ways.

Efficiency through knowledge reuse: Semantic knowledge graphs allow AI systems to leverage structured knowledge and perform reasoning without brute-force computation. By encoding facts and relationships explicitly, a knowledge graph can help an AI answer questions or make inferences by traversing graph connections, which is often far less computationally intensive than deep neural processing over massive text corpora. As one analysis notes, **knowledge-driven AI agents rely on structured data and rules and “require far less brute-force” computation compared to purely large language model approaches** ¹¹. In practice, Cruz's integration of knowledge graphs with AI (for example, in Retrieval-Augmented Generation setups) means the AI can be “grounded” with precise context ¹² ¹³. This reduces unnecessary guesswork and repeated trial-and-error processing on irrelevant data, thus potentially saving energy. By providing targeted, semantically relevant information to AI models (instead of huge volumes of unstructured text), graph-based context can cut down on the number of operations and the size of models needed to achieve a given accuracy ² ¹⁴. Over many queries or transactions, these efficiency gains translate into **lower energy consumption** and carbon output.

Optimization of IT workflows: Beyond AI model inference, Cruz's semantic approach can streamline general IT workflows. For example, consider the task of finding data or expertise within a large organization. Traditionally, consultants or engineers might perform multiple searches or run data-heavy analytics to locate relevant information. In a case study by Enterprise Knowledge, introducing a centralized sustainability **knowledge graph** (with a unified ontology for environmental and supply chain data) dramatically improved efficiency: consultants no longer had to dig through disparate documents and systems, saving time and computational resources in preparing sustainability insights ¹⁵ ¹⁶. The knowledge graph connected previously siloed data sources and provided a quick semantic search layer. As a result, the firm could more easily **advise on efficient measures to limit**

environmental impact, leveraging data-driven insights that were readily accessible instead of computed from scratch for each new project ¹⁷. This example illustrates how knowledge graphs act as a force-multiplier, reducing redundant data processing and enabling re-use of existing insights – an environmentally friendly outcome due to less wasted effort and compute.

Alignment with green IT initiatives: The focus on semantic data modeling also dovetails with industry efforts to make software and AI greener. For instance, **carbon-aware computing** often involves optimizing when and where computations run (e.g., scheduling tasks when renewable energy is available or using efficient algorithms). Knowledge graphs can assist by modeling energy data and system relationships. In fact, knowledge graphs have been applied in the energy domain to monitor consumption and suggest optimizations ¹⁸. While Cruz's primary applications are in cybersecurity and information management, the underlying principle – **model the problem domain explicitly to allow smarter, less wasteful computation** – aligns with green computing practices. A dynamic knowledge graph of an IT infrastructure, for example, could help route workloads to the most efficient resources, or automatically flag underutilized assets for consolidation. Such scenarios show the potential for Cruz's G³ philosophy (linking graphs across domains) to include environmental data graphs (e.g., data center power usage or carbon intensity of energy grids) linked with application and business logic graphs. By bridging these, an intelligent system could make holistic decisions that optimize for energy efficiency across the stack.

It's important to note that making AI truly sustainable is an ongoing challenge that extends beyond any single technique. However, **Cruz's semantic knowledge graph approach contributes by injecting structure and context into AI systems, thereby reducing reliance on raw computing power**. This complements other efforts like model compression, hardware efficiency, and renewable-powered data centers. Taken together, these strategies drive toward the ESG goal of minimizing the environmental footprint of IT. As AI thought leaders have emphasized, achieving sustainable AI means addressing both the supply side (greener AI development) and the demand side (AI for sustainability) ¹⁹. Cruz's work touches both: it aims to make AI development more efficient (supply side) and can be used to accelerate sustainability insights and data integration (demand side), as we will see in later sections.

Ethical AI and Inclusive Knowledge Frameworks

The **social** dimension of ESG in AI is about ensuring technologies are developed and used in ways that uphold ethical principles, human rights, and societal well-being ⁶. Key concerns include avoiding bias and discrimination, ensuring inclusivity, respecting privacy, and generally aligning AI with human values. Dinis Cruz's emphasis on semantic knowledge graphs and G³ inherently carries an **ethical AI orientation**, emphasizing transparency, diversity of perspective, and human-centric knowledge curation.

Avoiding one-dimensional bias: Traditional AI systems often suffer from biases in training data – if the data is skewed or reflective of historical prejudices, the AI's decisions will be too. One way to combat this is to introduce **explicit knowledge and rules** that can counteract or contextualize what pure data-driven learning provides. By designing **ontologies and knowledge graphs deliberately**, we can embed ethical constraints and a more balanced worldview into the system's knowledge base. Cruz's G³ concept – connecting multiple ontologies – is particularly relevant here. It recognizes that **no single taxonomy can capture all human perspectives**, and that attempting to enforce one "master ontology" could impose the biases or blind spots of its creators ³. Instead, G³ allows for multiple domain ontologies (even multiple cultural or value systems) to co-exist and be interlinked. This pluralistic approach is aligned with **AI ethics principles of inclusivity and fairness**, as it encourages the representation of diverse viewpoints. In a practical sense, an AI that consults a Graph-of-Graphs could cross-check

information: for example, a health AI might draw from both a medical knowledge graph and a patient community graph, ensuring it considers both clinical facts and patient perspectives. Such multi-graph reasoning can mitigate the risk of tunnel vision or biased recommendations that might arise from a single-source system.

Transparency in meaning-making: Ethical AI also demands **transparency and explainability** – stakeholders should be able to understand why an AI made a decision. Cruz has argued that interoperability isn't just about data exchange, **"it's about trust – and trust begins with shared meaning."** In discussions on semantic systems, he and colleagues posed the idea that an **ontology should be a living social contract rather than a fixed map**, constantly refined through community critique ²⁰ ²¹. This philosophical stance has concrete implications: if our AI systems base their reasoning on knowledge graphs whose structure and content are openly scrutinizable, then the *meaning* used in AI decisions is laid bare. For instance, if an AI's decision about loan eligibility references a knowledge graph link like "employment status -> high credit risk" (just as an example), one can examine that link, question it, and update it if it's unjust or outdated. In Cruz's semantic OWASP initiative (applying graphs to security knowledge), a similar logic is applied – security requirements and best practices are made explicit nodes and links, so that anyone (developers, auditors, users) can trace *why* a certain practice is recommended ²² ²³. Translating this to AI ethics: a knowledge graph-driven AI can expose the **conceptual reasoning path** (via graph relationships) behind its outputs, making it easier to identify biases or errors in the knowledge. This stands in contrast to black-box neural networks that often entangle facts in millions of opaque parameters.

Harnessing AI for social good: Another aspect of ESG (Social) is using AI proactively for positive social impact. Knowledge graphs have a role to play here too, as they are being used in domains like healthcare, education, and public policy to organize information and support informed decision-making. Cruz's work, while centered on cybersecurity and information management, demonstrates how the approach can be generalized. For example, his project "Graph-Powered Legal Knowledge: An Open, Distributed, GenAI-Assisted Roadmap" suggests using graphs to democratize legal and regulatory knowledge ²⁴ ²⁵. By making complex legal frameworks navigable via graphs, such tools could empower more people to understand their rights and obligations, promoting social justice and compliance. Moreover, Cruz's collaboration with GenAI indicates an effort to ensure AI systems behave **deterministically with provenance** – in other words, AI that can cite its sources and reasoning ²⁶. This focus on provenance aligns with ethical AI guidelines calling for AI to be *accountable* and *explainable* in terms of where its information comes from (to avoid, say, misinformation or unjustified decisions). Indeed, Cruz even applied ChatGPT in a "Deep Research" mode to document how Niklas Luhmann's system of notes (a human knowledge system) parallels modern knowledge graphs ²⁷. This kind of reflective research underscores the importance of learning from human knowledge organization methods that respected context and meaning – something very relevant to making AI more **human-centric and ethical**.

In summary, the semantic knowledge graph and G³ approach naturally foster an **ethical, human-aligned AI design**. By making knowledge explicit, multifaceted, and transparent, Cruz's methods help AI developers address biases, include diverse perspectives, and provide explanations. These qualities answer directly to ESG social objectives such as **fairness, inclusivity, and respect for stakeholders**. They also complement frameworks like the EU's ethical AI guidelines and initiatives like the **Partnership on AI**, which emphasize involving stakeholders in AI system design and ensuring AI's outcomes are beneficial and just. As AI and semantic web technologies converge, Cruz's work exemplifies how we might avoid the pitfall of encoding rigid or biased worldviews into our intelligent systems. Instead, we treat knowledge models as evolving, scrutinizable social artifacts – an approach very much in line with treating AI development as a **shared societal endeavor**, not just a technical feat.

Stakeholder Transparency and Open-Source Collaboration

Transparency is a foundational element of **ESG governance** and is tightly linked with the **social license** to operate technology. Stakeholders – whether they are customers, employees, investors, or regulators – now expect visibility into how systems work, how decisions are made, and how risks are managed. Dinis Cruz’s practices strongly emphasize **open collaboration and transparency**, particularly evident in his commitment to open-source tooling and knowledge-sharing in community settings. These practices not only accelerate innovation but also build trust and accountability, aligning with ESG goals around stakeholder engagement and ethical governance.

Open-source ethos: Cruz has a long history in the open-source security community (for example, through OWASP), and he continues this ethos by open-sourcing the tools and platforms developed in his projects. In the **Semantic OWASP** initiative, one notable component is **MGraph-DB**, a serverless graph database tailored for integrating LLM outputs into a knowledge graph. Cruz’s team created MGraph-DB and deliberately **released it under an open license** ²⁸. This means anyone – including the OWASP community and the public – can inspect, use, and modify the database engine. The paper highlights that leveraging such **existing open-source tools** jump-starts the semantic graph ecosystem at low cost and risk ²⁹. Equally important, open-sourcing ensures that the knowledge infrastructure itself is transparent. Stakeholders are not beholden to a proprietary “black box” for storing and retrieving knowledge; instead they have a say in its development and can trust it more because they can see how it works under the hood. Cruz’s work also references **OWASP SBot**, a collection of security automation scripts that is “*completely open source*” and had already represented security data in graph form ³⁰. By building on open community projects, the development of the semantic knowledge base remains accountable to the community. In ESG terms, this is a model of **governance through openness** – decisions about the knowledge schema, updates, and usage policies can be discussed in the open, and a broad set of contributors can weigh in. This collaborative governance can mitigate risks of unilateral control or opaque changes that might harm stakeholders.

Transparency for stakeholders: One of Cruz’s compelling proposals is “**externalizing compliance knowledge**” via semantic graphs to foster transparency and trust in ecosystems of partners and customers ³¹ ³². In the context of cybersecurity compliance, he envisions organizations sharing parts of their compliance knowledge graph with third parties. For example, a company could provide regulators or clients with a queryable view of how its internal controls map to various security standards ³³. This idea mirrors what is happening in ESG reporting: companies are increasingly expected to share **ESG metrics** (carbon emissions, diversity stats, etc.) publicly. Cruz explicitly draws that parallel, noting “*companies are increasingly pressured to share ESG metrics*”, and suggests that similarly, security or compliance data could be shared in a structured way ³⁴. The benefit is that stakeholders gain **continuous, machine-readable assurance** rather than periodic, static reports ³⁵ ³⁶. Translated to ESG, one could imagine a future where an organization’s sustainability commitments (say, its carbon footprint or supply chain labor standards) are published as a knowledge graph accessible to investors and watchdogs. Such transparency would enable real-time tracking and comparison, holding companies accountable. Indeed, Cruz points out that if compliance information is made public and queryable, it **creates accountability**: anyone can monitor changes, and if a company quietly backtracks on a commitment, it would be readily noticed ³⁷. This is directly aligned with stakeholder expectations in ESG – stakeholders want consistent, verified information, not glossy reports. By using open knowledge graphs as “living reports,” organizations could strengthen trust with stakeholders and **demonstrate integrity** in a way that static disclosures cannot.

Collaboration and community governance: Open collaboration isn’t only about code; it’s also about **shared standards and ontologies**. Cruz’s G³ approach implies that different organizations or communities might maintain their own graphs but agree on ways to link them. For effective

transparency, some common vocabulary or mapping is needed. We see this in emerging ESG data standards – for instance, frameworks for carbon accounting or workforce diversity metrics that allow apples-to-apples comparisons across firms. If each company publishes ESG data in a graph but uses completely different terms, stakeholders face a new interoperability problem. Cruz's work on creating ontologies for security standards ³⁸ ³⁹ provides a template: he suggests that standard bodies (like an EU agency ENISA in security, or analogously GRI/SASB in sustainability) could maintain **official ontologies** for regulations or metrics, and companies then publish their data mapped to those shared schemas ³⁹. This approach is highly relevant to ESG collaboration. It would enable an ecosystem where data flows more freely and comparisons are accurate. Importantly, it would be **open-source** or open-data at the level of definitions – the ontologies themselves are a public good maintained transparently. Cruz even likens this to financial data sharing via XBRL (an open standard for financial reporting) ⁴⁰, hinting that compliance and ESG data could undergo a similar standardization for transparency's sake. We already see steps in this direction: for example, the EU's Open Data Portal and efforts to encode legislation in machine-readable form ⁴¹. Cruz's proposals reinforce that **open standards + open implementations = empowered stakeholders**.

In summary, the practices championed by Dinis Cruz around open-source development, knowledge sharing, and semantic transparency align closely with the **ESG ideal of stakeholder empowerment**. By making tools and data open, he helps shift power towards stakeholders who can inspect, verify, and contribute. **Stakeholder transparency** isn't just a checkbox; it becomes a built-in feature of the system – whether it's a security program or an AI application – through semantic graphs that can be exposed in controlled, auditable ways. This stands in positive contrast to many legacy systems where stakeholders must trust an organization's word or wait for annual reports. The combination of **open collaboration and live transparency** yields a more accountable and inclusive way of governing technology, very much in spirit with ESG governance principles and modern compliance trends.

Responsible Data Governance and Knowledge Graphs

As organizations handle ever-growing volumes of data – including personal, sensitive, or regulated data – **responsible data governance** has become paramount. This falls under the **Governance (G)** in ESG but also touches on privacy and social responsibility. It involves ensuring data is managed ethically, in compliance with laws (like GDPR or other data protection regulations), and with proper oversight and quality control. Semantic knowledge graphs, as leveraged in Cruz's work, offer powerful capabilities for data governance by making relationships and rules explicit in a machine-readable form. Here we explore how Cruz's approach contributes to robust data governance and aligns with responsible data practices.

Data lineage and accountability: One of the challenges in data governance is tracing where data comes from, how it's transformed, and who accesses it. Knowledge graphs inherently model relationships, which makes them ideal for capturing **data lineage** and dependencies. In a governance context, one can use a knowledge graph to link a data element (say a customer's email address) to all the systems where it resides, the processes that use it, and the policies that apply to it ⁴². This holistic map greatly simplifies answering questions like "Where is personal data stored?" or "Who has access to this data point?" Traditional siloed databases make such queries extremely difficult, whereas a knowledge graph can answer them with a straightforward traversal. Cruz's orientation towards connected knowledge aligns with this need: by transforming documentation and records into an interconnected graph, we gain **real-time visibility into the data ecosystem**. For example, if Dinis's team builds a semantic graph for an organization's security controls, that same graph could be extended to cover data assets, linking controls to specific datasets or applications. Then, if a new privacy regulation comes out, one could query the graph to find all data assets lacking a required control. This dynamic querying is far more efficient than manual audits. In fact, industry experts note that

knowledge graphs help organizations track data lineage and usage, simplifying compliance with regulations like GDPR ⁴³ . By storing metadata (data about data) as first-class nodes and edges – such as data owners, classification tags (PII, financial data, etc.), and permitted uses – a knowledge graph becomes a living data catalog that enforces governance rules.

Modeling policies and regulations: Responsible governance isn't just about knowing where data is; it's also about ensuring proper policies are applied. Here, semantic graphs shine by allowing encoding of **rules and constraints** alongside data. For instance, an **"access control"** policy can be represented as relationships between *roles*, *datasets*, and *permissions* in the graph ⁴⁴ . With the graph, one can then automatically check for policy violations (like an instance of a role having access to data it shouldn't). Similarly, privacy requirements can be modeled: consider GDPR's *right to be forgotten*. A knowledge graph could map each user's data across all systems, so that a deletion request triggers a graph query returning all records to erase ⁴⁴ . This ensures completeness in compliance. Cruz's method of linking graphs is pertinent here – for example, a **compliance ontology** (with concepts like consent, retention period, encryption status) could be linked with a company's data graph. The result is an integrated view where every piece of data either complies with the linked policy nodes or is flagged. Research has indeed proposed that **knowledge graphs serve as the backbone for data governance in AI**, embedding laws and policies right into the data fabric ⁴⁵ ⁴⁶ . Cruz's pursuit of deterministic outputs with provenance in AI is another form of governance: by capturing how an AI result was derived, including which data points were used, he provides the audit trail necessary for accountability. If a decision is challenged (say, an automated loan denial), a provenance-aware system can show the chain of data and rules that led to it – crucial for governance and regulatory compliance.

Security and risk management: Data governance also overlaps with cybersecurity (a domain of Cruz's expertise). Knowledge graphs can improve **risk management** by revealing hidden connections and aggregating risk indicators. Cruz's concept of a graph-based threat model, for instance, could link assets, vulnerabilities, controls, and threat actors. From a governance perspective, this means a Chief Risk Officer could query "show me all systems that lack multifactor authentication and handle personal data," combining security and privacy viewpoints – something that might otherwise require correlating separate inventories. Indeed, Cruz and collaborators have demonstrated knowledge graphs for threat modeling and supply chain risk ⁴⁷ ⁴⁸ . These allow organizations to **see the big picture** of their operational risk and compliance status on a continuous basis, rather than through disjointed spreadsheets. By making such graphs collaborative and shareable (with due access controls), organizations can also bring in external auditors or partners to validate their governance posture more efficiently. This level of transparency in internal controls is increasingly sought in ESG reporting (for the Governance part), where companies might disclose their cyber governance practices or data protection measures to investors. A semantic representation could back up those disclosures with concrete, queryable facts rather than just narrative assertions.

To sum up, **semantic knowledge graphs act as a foundational technology for responsible data governance**. They provide a unified, understandable map of an organization's data landscape, policies, and processes. Dinis Cruz's approach – treating documentation, standards, and data as an interlinked graph – directly supports this by breaking down silos and making the implicit explicit. This approach aligns with modern governance frameworks that call for **integrated GRC (Governance, Risk, Compliance)** solutions and continuous controls monitoring. By comparing Cruz's graph-centric vision with external initiatives, we see harmony: for example, the **Basel Committee** has discussed principles for effective risk data aggregation, which emphasize completeness, accuracy, timeliness – all of which a well-constructed knowledge graph can enhance. Additionally, industry solutions are emerging that use knowledge graphs to ensure **regulatory compliance** automatically ⁴⁹ ⁵⁰ , validating Cruz's direction. Responsible data governance is ultimately about trust – ensuring stakeholders (from customers to

regulators) can trust that data is handled correctly. Semantic graphs, by offering **clarity, traceability, and enforceability** of data relationships, are powerful enablers of that trust.

AI Explainability and Accountability through Semantic Graphs

One of the most critical challenges in modern AI, especially complex machine learning models, is **explainability** – the ability to understand and trace how an AI system arrives at a given output. Explainability is tightly linked to accountability: if we can explain an AI decision, we can better hold the system (or its operators) accountable for its outcomes, and we can debug or improve the system as needed. This is reflected in ESG under both Social (e.g., avoiding harm through opaque decisions) and Governance (e.g., oversight of AI systems). Dinis Cruz's integration of semantic knowledge graphs into AI workflows provides a promising path to greater explainability, as evidenced by his focus on provenance and deterministic generation of content.

Provenance-enabled AI: In Cruz's experiments with Generative AI (GenAI), he advocates for **deterministic outputs with provenance** – meaning that an AI's answers should be repeatable and accompanied by source references ⁵¹. By grounding AI responses in a knowledge graph or a documented knowledge source, we can achieve a form of **AI that cites its work**. For instance, in a generative Q&A system enhanced by a knowledge graph, the answer to a user's query isn't just a free-form text; it can include pointers to the exact nodes or facts from the graph that informed the answer. This is already seen in Retrieval-Augmented Generation setups where source documents are cited. Knowledge graphs take it further by citing **fact nodes and relationships**, which are more granular and often easier to interpret than raw text excerpts. As one industry example highlighted, when a knowledge graph is used beneath an LLM, *"once we have a knowledge graph, we can walk back from our result to the information used to generate the result,"* providing **concrete and verifiable information** to users ⁵². The ability to walk back through the graph means a user (or auditor) can see the chain of reasoning: from the final answer, you trace to key intermediary concepts, and ultimately to source data or documents linked in the graph. This is a clear boon for explainability, effectively creating an **explanation graph** for each output on demand.

Semantic context as explanation: Another angle is that knowledge graphs provide **human-intelligible context** by nature. While a neural network's latent features are inscrutable, a knowledge graph deals in named entities and relations that make sense to people (e.g., *User -> located in -> Country; Country -> has law -> Data Privacy Act*). When an AI's internal reasoning is structured around such connections (as would be the case if it queries a graph for relevant knowledge), the resulting explanation can be communicated in those terms. For example, consider an AI system that flags a transaction as potentially fraudulent. If it uses a knowledge graph of fraud patterns, the system might explain: *"This transaction shares attributes with known fraud cases (same device ID and unusual time of day)"*, linking to those attributes in the graph. Contrast this with a black-box explanation like *"score 0.9 on neuron 512 exceeded threshold"* – the latter is meaningless to stakeholders. ESG principles and emerging regulations (like the EU AI Act) are increasingly demanding that **AI decisions be explainable in understandable terms** ⁸. By using semantic layers, Cruz's approach naturally produces **concept-level explanations**. Indeed, Cruz's writing on news and fact-checking underscores the need for *fact provenance* to combat misinformation ⁵³. If an AI summarizes a news article, a provenance-aware, graph-backed system could not only cite the article, but also indicate which statements align with verified facts in a knowledge network of previous news. This level of explainability builds trust with end-users – they can verify for themselves or at least see that the answer was constructed from known building blocks.

Comparative accountability: It's also insightful to compare Cruz's approach with other explainable AI initiatives. Approaches like LIME or SHAP (popular XAI techniques) often give *local explanations* for model predictions (e.g., highlighting which input features influenced a decision). Those are useful, but they don't provide broader context or guarantee consistency in explanations. A knowledge graph-backed explanation, in contrast, ties the decision to an external knowledge base that can be the same for everyone and consistently updated. This fosters **accountability** because if there's an error or bias in the explanation, it likely traces back to a specific part of the knowledge graph or rule that can be debated and corrected openly. It also means explanations are not just post-hoc rationalizations; they are part of the system's **actual reasoning process**. This is critical – explanations derived from how the system actually works (as in Cruz's deterministic, provenance-based approach) are more trustworthy than ones generated after the fact. As AI expert Joanna Bryson famously noted, *"AI should be as transparent as possible – it's easier to trust a system that can show you why it did something."* Cruz's work operationalizes this by combining AI with graphs to *show its work*. For companies concerned about ESG, having such traceable AI is not just about trust – it will likely be a compliance requirement. Frameworks like **ISO/IEC 22989 on AI transparency** or the proposed EU AI Act transparency provisions explicitly call for information on how AI decisions are made, especially for high-stakes uses ⁵⁴ ⁵⁵. By building on knowledge graphs, organizations can create a **documentation trail for AI decisions** automatically, aiding both internal audits and external compliance checks.

In essence, Dinis Cruz's melding of semantic knowledge techniques with AI responds directly to the cry for **explainable and accountable AI**. It offers a path to move beyond treating AI as an inscrutable oracle. Instead, AI becomes more of a **collaborative system**, where the knowledge graph component represents what AI *knows* (and this part is visible and understandable), while the machine learning component provides pattern recognition and generalization power. The result is a hybrid that can answer *and* explain. This aligns with the visions of organizations like **DARPA's XAI program** and global think-tanks which have been advocating for AI systems that users – be it doctors, judges, or consumers – can question and get sensible answers from. Ultimately, integrating explainability through semantic graphs supports the **Governance** aspect of ESG by enabling oversight and remediation, and supports the **Social** aspect by respecting the rights of those affected by AI decisions to understand and challenge outcomes. As AI continues to permeate decision-making, such capabilities will be indispensable for any organization claiming its AI is responsible and aligned with societal values.

Alignment with Broader Sustainable AI Initiatives

To put Dinis Cruz's contributions in context, it's helpful to compare and contrast with external initiatives aimed at making AI and IT more sustainable and aligned with ESG. The encouraging finding is that Cruz's focus on knowledge graphs, transparency, and open collaboration strongly complements many of the directions these initiatives are heading. We highlight a few key efforts and how they map onto the same landscape:

- **Green Software Foundation (GSF):** This industry consortium focuses on reducing software-related carbon emissions. In their position papers, they emphasize lifecycle accountability for AI models and encourage standards for measuring AI's environmental impact ⁵⁶ ⁹. Cruz's work doesn't directly measure carbon, but by advocating efficient knowledge-driven AI and serverless graph databases (e.g., MGraph-DB running cheaply in the cloud) ⁵⁷, he addresses the *how* of building low-footprint AI services. Both GSF and Cruz stress **collaboration** – GSF via open standards, Cruz via open-source tools – as necessary to achieve sustainability at scale. It's plausible that knowledge graph techniques could feed into GSF's toolkits (for example, a graph that tracks carbon costs of different model pipelines to help developers choose greener options).

- **EU AI Act and Regulatory Frameworks:** The EU AI Act (currently in progress) is notable for weaving in ESG-like objectives: not only requiring that AI be technically safe, but also **transparent and environmentally sustainable** where possible ⁸. It calls for documentation of training data (for transparency) and oversight of high-risk systems. Cruz's insistence on provenance documentation and knowledge capture directly supports compliance with such regulations. If an organization adopts Cruz's semantic approach, they would more easily meet requirements to explain AI decisions or to show that they avoid using data in prohibited ways. Moreover, Cruz's idea of sharing compliance graphs externally anticipates regulatory moves towards continuous auditing. In a world where regulators might demand *API access* to certain AI governance data, having a knowledge graph backend as he proposes would be an elegant solution ³³ ⁵⁸. Thus, Cruz's work is forward-compatible with stricter future rules that link ESG outcomes (like climate impact or fairness) to AI governance.
- **Partnership on AI and Ethical AI Frameworks:** The Partnership on AI (PAI) is a multi-stakeholder body that issues best practices on topics like fairness, explainability, and AI's societal impacts. One of their themes is **responsible data sharing and benchmark transparency**. Cruz's open knowledge graph approach could serve as a concrete implementation of PAI's high-level principles – for instance, enabling researchers and auditors to inspect the knowledge an AI was trained on (through the graph) echoes PAI's calls for transparency. Other frameworks like **OECD AI Principles** and **UNESCO's Recommendation on AI Ethics** emphasize inclusiveness, multi-stakeholder governance, and data governance – all values reflected in Cruz's G³ philosophy and open working methods. For example, UNESCO calls for "public engagement and awareness" in AI governance; sharing semantic knowledge graphs of how AI systems operate could be a powerful means of public engagement, turning opaque systems into understandable knowledge networks.
- **Sustainable AI for Climate and Social Good:** There is a growing movement of using AI itself to advance sustainability (often termed "AI for Good" or specifically AI for climate action, etc.). Knowledge graphs are recognized in these circles as critical for integrating sustainability data. Academic and industry projects have used knowledge graphs to link and analyze ESG datasets – for instance, to unify various ESG indicators and standards for better investment decisions ⁵⁹, or to improve supply chain transparency. Cruz's work on linking graphs across domains suggests that AI systems can draw on both internal data and external sustainability data. Imagine a corporate digital assistant using an internal policy graph and an external climate impact graph to advise on a decision (e.g., which supplier has the lower carbon footprint given equivalent cost). Indeed, one LinkedIn piece by graph expert Philip Rathle described combining LLMs with knowledge graphs specifically to revolutionize ESG research ⁶⁰. The **interdisciplinary integration** that Cruz advocates via G³ is exactly what's needed to connect, say, climate science data, with business process data, with social impact knowledge, into one reasoning loop for AI.
- **Open-Source and Open Data Initiatives:** Finally, open-source communities (like the Linux Foundation's TODO Group for open OSPO practices, or data.world's open data catalogs) indirectly support ESG by fostering transparency and shared progress. Cruz's contributions to open source (e.g., OWASP projects, open research publications) resonate here. Notably, his "ChatGPT Deep Research" collaborations (like the Luhmann briefing or this very white paper) demonstrate how AI and humans together can produce open knowledge resources on complex topics ²⁷. This is an inspiring model for how we might tackle ESG challenges: a mix of human insight, AI assistance, and open dissemination of findings. It aligns with the broader "open AI" ethos that solutions to societal challenges should not be locked behind proprietary barriers. By crediting ChatGPT as a co-author on research, Cruz is also pushing the envelope on **AI**

transparency in authorship, a small but symbolically important part of governance (acknowledging AI's role rather than obscuring it).

In conclusion, the landscape of ESG and sustainable AI initiatives is vast, but Dinis Cruz's work on semantic knowledge graphs and G³ finds a comfortable and meaningful place within it. Whether the focus is reducing carbon emissions of AI, ensuring ethical and explainable AI, or improving transparency and stakeholder trust, we see that **knowledge-centric approaches are increasingly essential**. Cruz's contributions exemplify this trend and even anticipate future needs (like dynamic compliance sharing or multi-ontology integration) that not all initiatives have fully addressed yet. The comparison suggests a convergence: *sustainable, ESG-aligned AI requires robust knowledge management and transparency*, exactly what semantic graphs provide. As organizations and consortia continue to develop standards and tools for ESG in tech, incorporating the kind of semantic techniques championed by Cruz will likely accelerate progress. Conversely, Cruz's efforts gain validation and broader context through these initiatives, highlighting that his specific projects are part of a much larger movement to align technology with humanity's sustainability and ethical imperatives.

Actionable Takeaways

In light of the above exploration, here are key takeaways and recommendations for practitioners, organizations, and stakeholders aiming to build sustainable and ESG-aligned IT/AI systems:

- **Leverage Knowledge Graphs for ESG Data Integration:** Break down data silos by constructing semantic knowledge graphs that unify information from diverse domains (environmental data, compliance records, social impact metrics, etc.). This will facilitate holistic analysis and reporting of ESG performance, as knowledge graphs *"enable semantic modeling and integration of data from different sources, particularly useful for the collection and processing of ESG metrics"* ⁶¹. Start with a pilot graph connecting two high-priority domains (e.g., link your enterprise asset database with a carbon footprint dataset) to demonstrate value.
- **Optimize AI Systems with Contextual Knowledge:** Incorporate domain knowledge via graphs to reduce computational waste and improve accuracy. Instead of relying solely on large models, use a knowledge graph to provide precise context or logic for the AI. This hybrid approach can **cut energy usage** by avoiding brute-force processing, aligning with carbon-efficiency goals. Track metrics like reduction in API calls or model size when a knowledge graph is introduced, and translate that into estimated energy saved or CO₂ reduced to quantify the environmental benefit.
- **Embed Ethical Constraints and Diversity in Knowledge Models:** When designing ontologies and graphs, involve a diverse group of experts and stakeholders to capture multiple perspectives. Treat the ontology as a living document that can be updated as values or regulations evolve. Implement review processes for knowledge graph content to identify and mitigate biases. For example, you might have an ethics review board examine new rules added to a decision-logic graph. This ensures the AI's knowledge base remains inclusive and fair, echoing Cruz's G³ principle of multiple viewpoints over one master view ³.
- **Adopt Open-Source Tools and Collaborate:** Utilize and contribute to open-source graph tooling and ESG data standards. Tools like Cruz's **MGraph-DB** (or other RDF/graph databases) are available under open licenses ²⁸ – piloting them can save cost and build community knowledge. Participate in industry groups or open standard initiatives (e.g., W3C's ontology standards for schema.org, or open sustainability data formats) to stay aligned with best practices. Open-

source your non-sensitive ESG-related code or ontologies to foster an ecosystem of transparency and collective improvement.

- **Implement Provenance and Explainability Features:** Ensure that your AI systems can “**show their work.**” Technically, this can mean logging not just decisions but the decision path: what data was consulted, which graph nodes were traversed, which rules fired. If using a knowledge graph, attach metadata for sources (citations, timestamps, authorship). This will make internal audits and external disclosures much easier. Users and regulators are increasingly demanding such explainability, and building it in by design will keep you ahead of compliance needs. A practical step is to use frameworks that support explanation graphs or to integrate a mechanism where every AI output comes with a provenance trace (as Cruz demonstrates in RAG + graph scenarios ⁵²).
- **Share ESG Knowledge Graphs with Stakeholders:** Consider publishing parts of your knowledge graphs to enhance stakeholder trust. For instance, an “**ESG commitments**” graph could be made accessible to investors or the public, showing how various initiatives (carbon reduction projects, DEI programs, etc.) link to outcomes and metrics. Even if full data sharing isn’t feasible, share the ontology and structure – this transparency signals accountability. Cruz’s concept of providing partners a view into compliance graphs ³³ can be emulated for ESG: provide your sustainability partners or auditors with API access to certain data (with appropriate safeguards). This not only builds trust but might streamline reporting and partnerships as well.
- **Use Semantic Tech to Automate Governance:** Deploy knowledge graphs to automate checks for policy compliance and risk. For example, encode regulatory requirements (GDPR, ISO 27001, environmental regulations) in the graph and run periodic queries to find violations or gaps. This can function as an early warning system for governance issues. It transforms governance from a periodic manual task to a continuous, intelligent oversight mechanism, aligning with the “**continuous compliance**” vision Cruz describes ³³ . It will also prepare your organization for real-time compliance reporting, a likely future requirement.
- **Measure and Communicate the Impact:** Finally, establish metrics to track how these practices improve sustainability and ESG outcomes. This could include energy saved (via efficiency gains), faster compliance reporting times, reduction in incidents (security or ethical), or improved stakeholder sentiment. Use the knowledge graph itself to store and relate these metrics to interventions. Then **close the loop by communicating these wins** in sustainability reports or governance disclosures, crediting the semantic approach. For instance, if adopting an open knowledge graph reduced the time to answer client ESG inquiries from weeks to days, include that as a case study in your ESG report – it demonstrates both operational excellence and a commitment to transparency.

By taking these actions, organizations will not only align with ESG objectives in principle but also in practice, embedding sustainability and ethics into the very architecture of their IT and AI systems. The work of Dinis Cruz exemplifies how cutting-edge research can be applied in practical ways to achieve these goals. The path forward is one of **integration** – integrating knowledge for smarter AI, integrating ESG values into technology, and integrating efforts across communities via open collaboration. This synergy of semantic technology and ESG strategy paves the way for IT systems that are not only intelligent, but also **responsible, resilient, and worthy of stakeholders’ trust.**

Conclusion

The convergence of Dinis Cruz's semantic knowledge graph approach with sustainable and ESG-aligned technology practices offers a compelling vision for the future of responsible innovation. By structuring knowledge as interconnected graphs (and even graphs of graphs), Cruz's work provides a blueprint for tackling complexity in a way that keeps systems **adaptive, transparent, and efficient**. These qualities are exactly what organizations need as they strive to reduce their carbon footprint, ensure ethical AI behavior, engage stakeholders with honesty, govern data responsibly, and demand explainability from automated decisions.

In this white paper, we have mapped Cruz's research and engineering practices to the key pillars of ESG:

- On the environmental front, semantic graphs and knowledge-driven AI can help bend the curve of carbon emissions, enabling more with less energy and integrating crucial sustainability data for smarter decisions.
- On the social and ethical front, the embrace of multiple ontologies and open collaboration embeds inclusivity and fairness into the fabric of AI systems, while enhancing trust through transparency and provenance.
- On governance, the fusion of knowledge graphs with compliance and policy management creates a robust framework for accountability – one that regulators, partners, and the public can have confidence in, because it's visible and verifiable.

Crucially, this alignment is not happening in isolation. It resonates with and reinforces broader initiatives – from industry consortiums like the Green Software Foundation to policy efforts like the EU AI Act and beyond. The comparison shows a common trajectory towards AI and IT that are **accountable to society and the planet**. What Cruz's perspective adds is a recognition that *better knowledge management is a cornerstone of this journey*. In a sense, to make AI sustainable and ethical, we must also "sustain" and curate the knowledge that AI uses. Graphs, ontologies, and open knowledge ecosystems turn out to be indispensable tools in doing so.

As organizations and practitioners, the opportunity now is to put these ideas into action. By building systems that mirror the interconnectedness of the real world and the values we aim to uphold, we can avoid the pitfalls of opaque, one-size-fits-all solutions. Instead, we get systems that are **context-aware, critique-friendly, and capable of continuous improvement** – qualities that define not just sustainable technology, but sustainable progress.

In closing, the work of Dinis Cruz exemplifies a path forward where advanced technology and ESG ideals are not at odds but in harmony. It's a path where *knowledge* is elevated to the same stature as data and code in our engineering priorities. Such a path leads to AI and IT systems that earn trust not by assertion, but by design – through **open knowledge, clear reasoning, and aligned purpose**. This is the kind of future-ready foundation that can help businesses, governments, and communities navigate the complex challenges of our time, from climate change to digital ethics, with confidence and integrity. The graphs we build today, imbued with meaning and responsibility, can become the guiding maps for a sustainable tomorrow.

Sources:

1. Cruz, D. (2025). *Why I focus on G3: Ontologies and Taxonomies of Ontologies*. LinkedIn post highlighting the need to connect multiple knowledge graphs without a single master ontology

2. Steuer, T. (2025). *Graphs for LLMs: A Visualization*. Discusses knowledge graphs as the next evolution for reliable AI, noting up to 54% accuracy improvements when supplying context to LLMs ⁶² ⁶³ .
3. KPMG (2023). *Decoding Sustainable AI vs. AI for Sustainability*. Defines sustainable AI from an ESG perspective, emphasizing reducing AI's carbon footprint, ethical AI practices, and robust governance ⁵ ⁶ ⁷ .
4. Strubell et al. (2019) via KPMG – Noted that training a medium NLP model can emit ~626,000 lbs CO₂, equivalent to 5 cars' lifetime emissions ¹⁰ , underscoring the urgency of carbon-efficient AI.
5. Enterprise Knowledge (2022). *Knowledge Graph for ESG Case Study*. Describes how a consulting firm used a semantic knowledge graph to unify sustainability data, saving time and enabling insights to **limit environmental impact** ¹⁵ ¹⁷ .
6. Milvus/Zilliz (2023). *Knowledge Graphs in Data Governance*. Explains that knowledge graphs simplify compliance by linking data to policies and enabling queries for GDPR-related tasks (e.g., data deletion requests) ⁴⁴ .
7. Cruz, D. (2025). *Semantic OWASP (Research Hub)*. Proposes a semantic knowledge graph for open OWASP content; notes **MGraph-DB is open-source** and serverless, easing adoption ²⁸ , and highlights leveraging prior open-source efforts (OWASP SBot) ³⁰ .
8. Cruz, D. (2025). *Scaling Europe's Regulatory Superpower (Research Hub)*. Suggests sharing compliance knowledge graphs with partners/regulators for continuous assurance, drawing analogy to ESG transparency ³³ ³⁴ . Emphasizes how transparency via graphs can improve trust and accountability in ecosystems ³⁶ ³⁷ .
9. European Parliament (2025). *EU AI Act Highlights*. States Parliament's stance that AI in EU should be **safe, transparent, and environmentally friendly** ⁸ , reflecting ESG priorities in policy.
10. Green Software Foundation (2023). *Responsible AI is Green AI*. Affirms that responsible AI must consider carbon emissions along with social impact ⁹ , advocating collaboration on standards for greener AI.
11. Rathle, P. (2023). *ESG analysis with GraphRAG*. (LinkedIn) – Suggests integrating LLMs and knowledge graphs can revolutionize ESG research ⁶⁰ , aligning with Cruz's approach of combining GenAI with structured knowledge for complex analysis.
12. EnergyCentral (2024). *Building Better Generative AI with Graphs*. Describes how adding knowledge graphs to a RAG chatbot improved accuracy and enabled tracing answers back to source data for verification ⁵² . Also references Ontotext's Transparency Energy Knowledge Graph as a real-world example of graph-driven data transparency in the energy sector ⁶⁴ .

¹ ² ¹⁴ ⁶² ⁶³ nice visualization and content on why graphs are the way to provide context to LLMs | Dinis Cruz

https://www.linkedin.com/posts/discruz_nice-visualization-and-content-on-why-graphs-activity-7333497516329484289-fbbsg

- 3 20 21 This is exactly why my focus is on G3 (Graphs of Graphs of Graphs), where we have ontologies of Ontologies and Taxonomies of Taxonomies | Dinis Cruz
https://www.linkedin.com/posts/diniscruz_this-is-exactly-why-my-focus-is-on-g3-graphs-activity-7336340555905318914-DiBs
- 4 5 6 7 10 19 Decoding Sustainable AI vs. AI for Sustainability - KPMG Netherlands
<https://kpmg.com/nl/en/home/insights/2023/10/decoding-sustainable-ai-vs-ai-for-sustainability.html>
- 8 54 55 EU AI Act: first regulation on artificial intelligence | Topics | European Parliament
<https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- 9 Responsible AI is green AI - State of Green Software
<https://stateof.greensoftware.foundation/en/insights/responsible-ai-is-green-ai/>
- 11 Energy Efficiency of AI: LLMs vs Agentic Knowledge Graphs - LinkedIn
<https://www.linkedin.com/pulse/energy-efficiency-ai-llms-vs-agentic-knowledge-graphs-rick-gillespie-kwiye>
- 12 13 52 64 Building a Better Generative AI—How Knowledge Graphs Will Improve Our RAG-Powered Chatbot
<https://www.energycentral.com/home/post/building-better-generative-ai--how-knowledge-graphs-will-improve-our-rag-PmqhK49gbanPuWe>
- 15 16 17 Knowledge Graph Accelerator for Data Standardization and Environmental Social Governance (ESG) - Enterprise Knowledge
<https://enterprise-knowledge.com/knowledge-graph-accelerator-for-data-standardization-and-environmental-social-governance-esg/>
- 18 Energy Knowledge Graphs → Term
<https://energy.sustainability-directory.com/term/energy-knowledge-graphs/>
- 22 23 28 29 30 57 Semantic OWASP: Leveraging GenAI and Graphs to Customise and Scale Security Knowledge - Dinis Cruz - Research Hub
http://docs.diniscruz.ai/2025/04/02/semantic-owasp__leveraging-genai-and-graphs-to-customise-and-scale-security-knowledge.html
- 24 25 47 48 51 53 Dinis Cruz - Documents and Research
<http://docs.diniscruz.ai>
- 26 27 PDF: Bridging Niklas Luhmann's Ideas with Semantic Knowledge Graphs and G³ | Dinis Cruz
https://www.linkedin.com/posts/diniscruz_bridging-niklas-luhmanns-ideas-with-semantic-activity-7341126921713319938-HIbR
- 31 32 33 34 35 36 37 38 39 40 41 58 Scaling Europe's Regulatory Superpower: From Static Cybersecurity Standards to Semantic Graphs - Dinis Cruz - Documents and Research
<http://docs.diniscruz.ai/2025/03/31/scaling-europe-regulatory-superpower.html>
- 42 44 49 How do knowledge graphs help in data governance?
<https://milvus.io/ai-quick-reference/how-do-knowledge-graphs-help-in-data-governance>
- 43 Knowledge Graphs: The Key to Modern Data Governance
<https://www.actian.com/blog/data-governance/knowledge-graphs-the-key-to-modern-data-governance/>
- 45 KG.GOV: Knowledge graphs as the backbone of data governance in AI
<https://www.sciencedirect.com/science/article/pii/S1570826824000337>
- 46 Leveraging Knowledge Graphs for Governance in the AI Fabric - Altair
<https://altair.com/blog/articles/leveraging-knowledge-graphs-for-governance-in-the-ai-fabric>
- 50 Compliance Reasoning with Knowledge Graphs - Relational AI
<https://relational.ai/resources/compliance-reasoning-with-knowledge-graphs>

56 Green AI Position Paper | GSF - Green Software Foundation

<https://greensoftware.foundation/articles/green-ai-position-paper/>

59 [PDF] Shared Metrics of Sustainability: a Knowledge Graph Approach

<https://ceur-ws.org/Vol-3194/paper30.pdf>

60 ESG analysis, brought to you by GraphRAG. | Philip Rathle - LinkedIn

https://www.linkedin.com/posts/prathle_leveraging-large-language-models-knowledge-activity-7200501682894508032-mwe_

61 [PDF] Leveraging knowledge graphs for ESG metrics extraction

<https://dl.gi.de/bitstreams/35525b4b-d33a-4282-a7a6-5d79c81b1a76/download>