



Where to focus first on battery decarbonisation

Dr Robert Pell, Founder & CEO, Minviro



Disclaimer

Materials by Minviro

The services and materials provided by Minviro Ltd are subjected to Minviro's Standard Terms (a copy of which is available upon request) or such other agreement as may have been previously executed by Minviro Ltd.

The materials contained in this presentation are designed for the sole use by the presenter and solely for the limited purposes described in the presentation. The material shall not be copied or given to any person or entity other than the participants in this presentation without the prior written consent of Minviro Ltd.

The content presented in this presentation should be used with careful interpretation and consideration. Minviro shall have no liability whatsoever to any Third Party with regard to the use of materials provided in this presentation, including the accuracy or completeness of there of.

Our Expertise

Delivering LCA for enterprise supply chains since 2018

BATTERY & EV MANUFACTURERS

TESLA

STELLANTIS

SUNWODA
ENERGY



JLR

VERKOR

Gotion

MINING & RAW MATERIALS

RioTinto

CMOC

Pilbara
Minerals



NOVONIX

ExxonMobil

III
III
III
SOUTH 32

BATTERY MATERIALS & CHEMICALS

BASF

one

talga

arcadium
lithium

WILLAMSBURY
ZERO CARBON LITHIUM™

PowerCo

SYRAH
RESOURCES

RECYCLING & CIRCULAR ECONOMY

RecycLiCo™
BATTERY MATERIALS

altilium
clean technology

ASCEND
ELEMENTS

Primobius
Battery recycling without acids

green
union

INDUSTRY ASSOCIATIONS AND OTHERS

TE TRANSPORT &
ENVIRONMENT

COBALT
INSTITUTE
building technology, advancing society

Nickel
INSTITUTE
Knowledge for a brighter future

Battery
Norway

VDA

G-BA BATTERY
PASSPORT


RECHARGE

Unique Ecosystems in XYCLE

Proven at Enterprise scale for cradle-to-grave
PCF & LCA under evolving regulations.


World's biggest *Electronics* company

Directly connecting data from
suppliers to ensure conformity
and **measure scope 3**,
all within XYCLE.




Orkla, *The Largest* FMCG in Scandinavia

Integrating operational data
within their data lake to **automate**
PCF calculations on their 10k
product portfolio.



Global automotive OEM

20,000+ BOM Line Items
BOM imported and structured
Migration to XYCLE from
alternative LCA softwares



Software Democratised LCA, AI is Industrialising It

01

Software Democratisation

Historically, software platforms successfully democratised Life Cycle Assessment (LCA).

This critical shift made LCA accessible far beyond traditional academic circles, laying the foundation for modern sustainability tracking.



02

AI Industrialisation

Today, Artificial Intelligence is actively industrialising the entire LCA process.

As a direct result of this shift, the financial and temporal costs associated with producing an LCA result are collapsing rapidly.



03

The Credibility Gap

A critical disparity has emerged: the effort to understand and verify results is not decreasing at the same rate.

This growing gap between generation and comprehension is where methodological credibility gets lost. We must restore balance to ensure robust, scientifically sound outputs.

THE AUTOMATION IMPERATIVE

Why This Matters Now

RESOURCE BURDEN

70%–80%

of total LCA time and cost is consumed by data collection and supply chain linking.

This disproportionate resource burden makes data collection the most obvious and high-value target for automation.

Source: Miah et al. (2018), cited in Popowicz et al. (2025)

Surging Market & Regulatory Demand

LCA demand is rising at an unprecedented rate, driven by explicit requirements:

- ◆ Digital Product Passports
- ◆ Battery Carbon Footprints
- ◆ Product EPDs
- ◆ Direct Supplier Requests

The Capacity Gap

Practitioner numbers are not rising fast enough to meet surging demand. Consequently, **AI will inevitably be deployed at scale**, regardless of whether the discipline is fully prepared for its implications.

Presentation Roadmap: Navigating AI in LCA

This seminar will systematically explore the integration of Artificial Intelligence into Life Cycle Assessment through five key areas:

01

AI Across the LCA Workflow

An examination of current AI capabilities and applications across standard LCA phases.

02

Demonstration of AI-Enabled Model Creation

A practical look at automated system definition, inventory generation, literature extraction, and gap filling.

03

Demonstration of AI-Assisted Human Workflows

Showcasing how AI can support practitioners in complex tasks like background dataset mapping.

04

Implications for the Discipline

An analysis of what this technological shift means for the future of LCA, the wider industry, and the day-to-day role of practitioners.

05

Contextual Application

Understanding why the appropriate use of AI differs fundamentally between ecodesign applications and strict regulatory reporting.

What the Literature Says: AI Across the ISO 14040/44 Workflow

A comprehensive analysis of five recent literature reviews reveals that AI is actively being applied across every single phase of the ISO 14040/44 framework.

PHASE 01

Goal and Scope

AI is utilised for estimating product lifetimes, defining temporal scopes, and suggesting optimal system boundaries.

PHASE 02

Inventory Analysis

Machine learning models are deployed for filling missing data, estimating unit processes, creating surrogate process models, selecting proxy datasets, and managing allocation procedures.

PHASE 03

Impact Assessment

Algorithms are increasingly predicting characterisation factors (with a strong emphasis on toxicity) and forecasting full LCA results directly.

PHASE 04

Interpretation

AI tools assist in hotspot identification, conducting sensitivity and uncertainty analyses, and scenario optimisation.

Key sources: Romeiko et al. (2024), Ghoroghi et al. (2022), Popowicz et al. (2025), Salla et al. (2025), and Nwagwu et al. (2025).

The Dominant Pattern: Predictive Machine Learning

Predictive modeling dominates published literature to estimate environmental impacts

Core Algorithmic Workhorses

Most published work utilizes machine learning to predict environmental impacts or inventory values based on specific product or process attributes.

Primary Models:

- Neural Networks
- Random Forests

As noted by Salla et al. (2025) and Romeiko et al. (2024).

High-Accuracy Performance

$$R^2 > 0.9$$

Performance metrics are highly promising in controlled environments, frequently exceeding 0.9 accuracy in building studies.

Ensemble Methods

Random Forest & Gradient Boosting deliver top-tier predictive capability, albeit at a higher computational cost.

Luan & Jing (2026)

Historical Context

2000–2001

This core concept is not entirely novel. The earliest examples of surrogate LCAs date back over two decades.

Initially developed for evaluating product concepts in early design stages to streamline sustainability metrics before full production.

Sousa et al.

The Small-Data Problem in AI-LCA

While predictive models show promise, they are severely constrained by pervasive data volume limits, compromising generalisation and industrial application. There are also significant issues with hallucination or over confidence.

< 1,500 (70%)

DATASET SAMPLE SIZE

Data Volume Deficit

According to **Romeiko et al. (2024)**, over **70%** of reviewed ML-LCA studies were trained on datasets with fewer than **1,500 samples**.

This severe lack of training data creates fundamental generalisation issues across predictive lifecycle models.

15% – 30%

ACCURACY LOSS

Typology Transfer Failure

Luan & Jing (2026) demonstrated that models trained on one specific building typology lose **15% to 30%** of their predictive accuracy when applied to a different typology without retraining.

This highlights poor model adaptability across real-world domains.

$R^2 < 0.75$

AT < 500 SAMPLES

Industrial Unreliability

Neural networks degrade sharply on small datasets. When trained on **<500 samples**, validation R^2 consistently drops below **0.75**.

Small-sample models lack the robustness required for reliable, scalable enterprise applications.

The Prototype-to-Practice Gap

Evaluating Academic Literature Findings vs. Real-World Deployment & Performance Metrics

Academic Literature Findings

Popowicz et al. (103 papers) & Luan & Jing (52 papers)

81 / 103

Case Illustrations

Theoretical examples only

18 / 103

Purely Conceptual

Frameworks without testing

38%

Simulation Only

Restricted to synthetic environments

4 / 103

Implemented

Partly or broadly applied in practice

Real-World Deployment & Reliability

Empirical Deployment Rates & Prediction Margin Analysis

8% **Real-World Deployment Rate**

Models reaching actual field deployment with mandatory post-occupancy verification (Luan & Jing review of 52 papers).

12%–25% **Prediction Error Margin**

Discrepancy margin where AI predictions missed actual outcomes in the rare instances real-world validation occurred.

HEADLINE CONCLUSION

The literature successfully proves the technical feasibility of AI in LCA, but it does not yet prove its reliability for practice.

Critical Gaps: What the Literature Under-Covers

The current body of research on AI in LCA has significant blind spots that undermine practical credibility.

Uncertainty Analysis

17% of studies include formal uncertainty analysis

Formal uncertainty analysis is alarmingly rare, present in only 17% of studies (Luan & Jing 2026).

Explainability

~12% utilize interpretable AI techniques

Techniques designed to make AI decisions interpretable, such as SHAP or LIME, are utilised in only about 12% of the published studies.

Dynamic Modeling

Strictly Static models predominate

The concept of dynamic, real-time LCA remains largely unexplored, with the vast majority of current models being strictly static (Salla et al. 2025).

Database Sensitivity

15% to 35% result swings

AI models are highly sensitive to data structures. Changing background databases (ecoinvent vs GaBi) caused results to swing 15% to 35% for identical scenarios (Luan & Jing 2026).

The Paradigm Shift

From Prediction to Workflow Automation

Literature Lag: While academic literature still focuses on "classic ML" (predicting environmental impacts with regressors), practice since 2024 has shifted toward LLMs and autonomous agents executing complex workflow tasks.

Classic ML

Predicting Impacts

- **Primary Role:** Uses regression models to estimate or predict final environmental impact metrics directly.
- **Scope:** Evaluates inputs/outputs numerically without engaging in system modeling.
- **Function:** Acts merely as a mathematical estimator producing a static final answer.

LLM Agents (Post-2024)

Building Models & Workflows

- **Execution Capabilities:** Reads literature, builds process flowsheets, maps datasets, and drafts comprehensive reports.
- **LCA Phase Mapping:** Formally mapped onto goal & scope, impact category selection, and interpretation (Nwagwu et al, SINTEF 2025).
- **Core Risk Flagged:** Hallucination is identified as the central risk in this automated paradigm.

Core Shift: AI is no longer merely predicting the final answer—it is actively building the underlying model that produces the answer.

Two Distinct Modes of AI in LCA



To understand the practical application of AI in LCA, we must distinguish between two primary operational modes:

	Mode 1: AI Builds the Model	Mode 2: AI Assists the Human
Function	The AI drafts the system boundaries, inventory, sources, and executes gap fills.	The AI speeds up specific, expert-level tasks, such as background dataset mapping.
Human Role	The practitioner acts as a rigorous Reviewer and Editor .	The practitioner remains the primary Decision-maker , equipped with a much faster tool.
Main Risk	The generation of a model that is structurally plausible but scientifically indefensible.	The human operator over-trusting the AI's automated suggestions.
Credibility Lever	Ensuring strict provenance is attached to every single number generated.	Complete transparency regarding exactly why a specific match or decision was suggested by the AI.

Demo 1 Setup: From Question to First-Draft Model

Mode 1: AI building the initial LCA model

1. Input Specification

The AI requires three core input parameters to formulate the draft model:

- Functional Unit
- Production Route
- Specific Goal

EXAMPLE INPUT QUERY

"1 kg battery-grade lithium carbonate from spodumene."

Defines the scope, baseline chemical target, and primary extraction pathway for automated generation.

2. AI Execution (4 Sequential Steps)

STEP 01 — System Boundary & Flowsheet

Proposes a comprehensive system boundary and a process flowsheet tailored to the defined goal.

STEP 02 — Literature Data Extraction

Searches and extracts relevant data from the literature, including yields, reagents, and energy requirements.

STEP 03 — Foreground Inventory Construction

Constructs the complete foreground inventory based on all synthesized and extracted literature data.

STEP 04 — Gap Filling & Estimation Flagging

Identifies and fills remaining data gaps, explicitly flagging these estimations for user review.

OBSERVATION BEAT 1

Transparent Provenance

Pay close attention during the demo indicating exactly where every number came from across all literature sources.

OBSERVATION BEAT 2

Methodological Boundaries

Note the specific methodological choices that the tool explicitly refuses to decide, forcing human expert intervention.

Post-Demo Recap: What Just Happened

The demonstration showcased a massive compression of effort: what traditionally requires hours or days of manual literature extraction was completed in mere minutes. However, it is vital to recognise that the output is strictly a draft model, not a finalised LCA result.

Crucially, three fundamental methodological judgements were intentionally left to the human practitioner:

01

System Boundary Definition

Deciding exactly what is included in the assessment, what is excluded, and the scientific justification for those boundaries.

02

Co-product Allocation

Determining the appropriate allocation method—whether mass, economic, or system expansion—based on the specific context.

03

Uncertainty Assessment

Evaluating how much the final result can be trusted and determining its suitability for the intended purpose.

Mandatory Rules for AI-Built Models

To maintain credibility, AI-generated LCA models must adhere to strict operational rules:

01. Absolute Provenance

Every single number must carry clear provenance, detailing its source, the extraction method used, and a definitive label stating whether the data is measured, derived from literature, estimated, or a proxy.

02. Transparent Gap Filling

AI-generated gap fills must be explicitly labelled. They must never be silently blended into primary data streams.

03. Explicit Methodological Choices

All methodological decisions must be surfaced explicitly to the user, rather than being buried within automated defaults.

04. Guaranteed Reproducibility


The model must be reproducible, meaning the exact same inputs should consistently yield the exact same model structure. This is a significant challenge, as standard LLMs do not guarantee reproducibility by default (a vulnerability rated highly by Nwagwu et al. 2025).

THE UNGLAMOROUS BOTTLENECK

Background Mapping in LCA

Every single foreground flow must be accurately linked to a corresponding background dataset—covering electricity, reagents, transport, and raw materials.

- **Highly Repetitive & Manual:** Heavily dependent on scarce expert knowledge and manually executed despite software advances.
- **Notoriously Error-Prone:** The specific phase where the dangerous concept of 'close enough' quietly compounds.
- **Significant Impact:** Leads to large cumulative errors in the final life cycle impact assessment.



73%

rely entirely on manual mapping

Even in advanced BIM-LCA integrations during impact assessment (Luan & Jing, 2026).

DEMO 2 SETUP

AI-Assisted Background Dataset Mapping

Mode 2 workflow: AI proposes candidate background datasets with explicit matching logic, while practitioner retains total control.

01 Ranked Suggestions

Unmapped Flow Analysis: AI automatically identifies missing LCI dataset connections across multi-tiered material models.

Candidate Pool: Generates a structured list of potential background datasets for each unmapped item.

Multi-Criteria Evaluation: Scores candidates based on geographical relevance, tech match, and temporal fit.

02 Explicit Reasoning

Matching Rationale: Provides clear narrative explanation detailing why specific candidates are ranked highest.

No-Match Behavior: Explicitly flags flows when no high-quality database match exists rather than assigning poor proxies.

Transparent Metrics: Displays exact fit scores across geography, technology, and time periods.

03 Human Control

Full Authority: Human practitioner retains power to accept, reject, or completely override AI recommendations.

Manual Overrides: Allows seamless selection of alternative datasets or custom proxy configurations.

Override Logging: Captures practitioner decisions and rationale for complete versioning and audit trails.

Design Principle: Why Assist, Not Automate

Background dataset mapping errors are completely invisible in final aggregated impact metrics.

High Risk

Full Automation

- **Invisible Mapping Errors:** An incorrect impact assessment looks identical to a correct one in aggregate metrics.
- **Promotes Complacency:** Makes expert reviewers lazier, encouraging blind acceptance of flawed underlying data.
- **False Efficiency:** Removes human oversight without providing any mechanism to catch systemic misallocations.

Recommended

Expert AI Assistance

- **Highlights Anomalies:** Proactively surfaces potential mapping anomalies, making edge-case errors easier to spot.
- **Faster Scrutiny:** Accelerates the expert reviewer's workflow while preserving critical human judgment.
- **Cost-Effective Rigor:** Reduces validation costs without discarding necessary quality control.

Core Principle: The technology must make rigorous scrutiny cheaper and more efficient—not render it unnecessary.

The Changing Role: What Practitioners Will Do Less Of

As AI tools become deeply integrated into the LCA workflow, day-to-day tasks will shift away from repetitive, low-value manual work.

01

Manual data entry and the tedious transcription of figures from academic papers and technical reports.

02

Endlessly searching through complex background databases to find the 'closest' dataset match.

03

Building standard, boilerplate models from scratch for routine assessments.

04

Formatting reports and drafting the initial, descriptive write-ups of the methodology and results.

05

Performing repetitive manual recalculations to accommodate minor scenario variants and sensitivity checks.

The Changing Role: What Practitioners Will Do Less Of

As AI tools become deeply integrated into the LCA workflow, day-to-day tasks will shift away from repetitive, low-value manual work.

01

Manual data entry and the tedious transcription of figures from academic papers and technical reports.

02

Endlessly searching through complex background databases to find the 'closest' dataset match.

03

Building standard, boilerplate models from scratch for routine assessments.

04

Formatting reports and drafting the initial, descriptive write-ups of the methodology and results.

05

Performing repetitive manual recalculations to accommodate minor scenario variants and sensitivity checks.

The Changing Role: What Practitioners Will Do More Of

Freed from manual data processing, the LCA practitioner's role will elevate to focus on high-value, strategic judgement.

Framing the Core Question

Ensuring the goal, scope, and functional unit accurately match the actual business or regulatory decision.

Methodological Calls

Making and defending complex methodological calls regarding system boundaries, allocation rules, and cut-off criteria.

Reviewing AI Models

Reviewing and verifying AI-generated models, acting more as a critical reviewer than a primary author.

Securing Primary Data

Securing high-quality primary data directly from suppliers—this remains the absolute biggest credibility lever and is the one thing AI cannot invent.

Communicating Nuance

Communicating nuance and uncertainty to decision-makers who typically demand a single, definitive number.

Portfolio-Scale Execution

Executing portfolio-scale assessments that were previously not economically viable.

Macro Implications for the LCA Industry

The industrialisation of LCA via AI carries profound structural implications across four key industry dimensions.

01

Volume vs. Expertise

The total volume of LCAs produced will grow far exceeding qualified experts. The most scarce assets shift from software access to trusted data and human judgement.

02

Value of Primary Data

Because AI can instantly generate an endless supply of plausible secondary numbers, primary data sourced directly from real supply chains becomes exponentially more valuable.

03

Verification Bottlenecks

Verification and critical review will emerge as primary industry bottlenecks, requiring the profession to find scalable methods to deliver effective oversight.

04

AI Environmental Footprint

The industry must account for AI's own environmental impact; heavy compute loads should be quantified and included in assessments where material (Nwagwu et al. 2025).

Not All LCAs Are the Same: The Use-Case Spectrum

The appropriate application of AI in LCA depends entirely on the specific use-case, which can be evaluated along a single critical axis: what are the consequences if the generated number is wrong?

LOW ERROR CONSEQUENCE

HIGH ERROR CONSEQUENCE

LOW CONSEQUENCE

Ecodesign & R&D Screening

Tolerance: High tolerance for error.

Requirements: Requires rapid speed and directional accuracy.

AI Role: AI can be heavily utilised for fast modeling.

MODERATE CONSEQUENCE

Internal Decisions & Procurement

Tolerance: Moderate consequence.

Requirements: Requires a balance of speed and improved accuracy.

AI Role: Clear understanding of data provenance required.

HIGH CONSEQUENCE

Reporting & Regulation

Tolerance: Extremely high consequence with low error tolerance.

Requirements: Absolute auditability, strict standards adherence, and human accountability.

AI Role: Restricted to transparent assistance.

Ecodesign: Where AI Earns its Keep Fastest

01

High Leverage & Low Data

Ecodesign is the domain where AI delivers immediate, high-impact value.

Early-stage design possesses the highest leverage for environmental improvement, yet suffers from the least available data.

This environment is exactly where rapid predictive modeling provides immense benefit (42 Technology, 2025).

02

Relative vs. Absolute Focus

Relative answers (determining that 'Design A is better than Design B') are significantly more robust to underlying data errors than absolute footprint calculations.

By embedding AI-driven impact feedback directly into engineering tools like CAD and PLM software, engineers can iterate rapidly before material decisions are permanently locked in.

03

Avoiding False Precision

The primary risk to monitor during early AI-driven ecodesign is false precision.

A rough screening number generated by AI must not be subsequently quoted or published as if it were a verified, absolute footprint.

Reporting and Regulation: Where Rigour is Non-Negotiable

Traceable, audit-ready data provenance vs. black-box optimization

Non-Negotiable Auditability

Regulators and verifiers demand a fully traceable audit trail. Every number, choice, and source must be open to scrutiny.

Black-box AI predictions inherently fail this mandatory verification process.

"Explainability & strict data provenance are mandatory entry requirements."

— Weerasinghe (2025), CSRD Audit Trails

Fixed ISO/PCR Standards

Established frameworks dictate rigid methodological choices that AI cannot optimize around:

- ISO 14044 & ISO 14067
- Product Category Rules (PCRs)
- Specific Battery Regulation rules

Methodological compliance is fixed and non-flexible.

Restricted AI Role

In high-stakes reporting, AI is strictly confined to supportive tasks:

- Accelerating data assembly
- Performing consistency checks
- Flagging data outliers
- Generating documentation

Human practitioner remains the solely accountable author.

A Simple Decision Rule for AI Deployment



To navigate the complexities of AI in LCA, practitioners should apply a simple, three-question decision rule before deploying AI tools:

1 Consequence

Who is relying on this final number, and what are the legal, financial, or reputational consequences if the number is wrong?

2 Traceability

Can the system definitively show where every single data point and assumption came from?

3 Judgement

Is the task at hand a purely data-driven operation, or does it require a methodological, scientific choice?

The Decision Rule

AUTOMATE

If it is a pure data task with low consequences for error, you can safely automate.

ASSIST

If it involves a methodological choice or carries high consequences, the AI must only assist, and the human must make the final decision.

Five Commitments for Credible AI-Enabled LCA

To ensure the future credibility of our discipline, we must collectively adopt these five commitments for AI-enabled LCA:

01. PROVENANCE

Provenance on Every Number

The source—whether measured, literature, estimated, or proxy—must always be explicitly visible.

02. HUMAN OWNERSHIP

Human Ownership of Methodology

Human practitioners must retain absolute ownership over system boundaries, allocation methods, and uncertainty assessments.

03. MATCH AUTOMATION

Match Automation to Consequence

AI deployment must be tailored to the use-case; an ecodesign screening tool is fundamentally different from an EPD generator.

04. REAL-WORLD VALIDATION

Validate Against Reality

AI models must be validated against actual, real-world physical outcomes, not merely benchmarked against other models.

05. PRIMARY DATA

Invest in Primary Data

The industry must relentlessly pursue high-quality primary data from supply chains, as this is the one critical asset that AI cannot generate.



Conclusion & References

"AI makes Life Cycle Assessment incredibly cheap to produce. Our collective job as a profession is to ensure it remains worth believing."

References

- Romeiko X.X. et al. (2024)
- Salla J.V.E. et al. (2025)
- Ghoroghi A. et al. (2022)
- Popowicz M. et al. (2025)
- Nwagwu C.C. et al. (2025)

- Luan K., Jing Y. (2026)
- Weerasinghe S. (2025)
- de Jesus J.O. et al. (2021)
- 42 Technology (2025)
- Sousa et al. (2000-2001)