Introducing TRACE

TRansport **AI** **C**arbon **E**valuator

Making the carbon impacts of AI visible in decision-making

CONTENTS

Executive Summary

The problem: Carbon impacts of AI aren't visible in decision-making

AI tools have the potential to drive carbon reductions further and faster in the transport sector in support of net zero, but this is set against the large and exponentially increasing carbon impacts of the hardware, training and use of AI.

AI should be seen neither as helpful nor harmful. Instead, decision makers must carefully weigh the different benefits and downsides of AI tools, case by case, thinking about how and where they are used in practice today.

As AI applications are rapidly rolled out at scale in all areas of the transport sector, decision-makers have no consistent, robust and timely means to measure a tool's carbon impacts and consider them alongside the applications' intended benefits.

This limits the ability of government bodies to realise AI's potential to support the transition to net zero and to undertake their statutory responsibility to give 'due consideration' to the environment.

The solution: Introducing TRACE

TRACE (**TR**ansport **AI** **C**arbon **E**valuator), developed by Mott MacDonald and City Science on behalf of the Department for Transport (DfT), is a **first-of-its-kind** tool that supports policy and decision-makers to a rapid assessment of the carbon impacts of AI applications, compared to a realistic counterfactual, and intended benefits. Analysis by the tool can be completed in just two to three hours and targets the early stages of development while there is still an opportunity to inform decisions on whether and how AI should be deployed.

Immediate results: Benefits to DfT and the Transport Sector

The DfT is looking to include TRACE as part of the its decision-making processes, with work underway to explore including a TRACE report within the AI Project Oversight Group's standard governance considerations.

Carbon impacts can now be consistently included in decision-making for AI tools, alongside value

Using TRACE Who should use it and how?

The tool is intended to support informed users in teams actively developing AI applications. Information on the AI model type, training, counterfactual and benefits are inputted, generating a clear output to inform decision-making processes and forums, including a detailed breakdown of key assumptions.

Development journey: Research to support this decision-useful tool

The tool is informed by a detailed literature review and carbon expertise and balances the need for a lightweight tool in the context of this fast-moving space, with the robustness and transparency in the calculation methodologies.

Where next for TRACE? Cross-sector and cross-government potential

Interviews undertaken during development of TRACE demonstrated that the lack of visibility of AI's carbon impacts in decision-making was present across the transport sector and beyond.

Real world testing and refinement is needed to be as impactful as possible. As TRACE is rolled-out within the DfT lessons learnt will be captured. In parallel, the opportunity to bring the tool to the wider transport sector, and other areas of government will be explored.

The problem:

Carbon impacts of AI aren't visible in decision-making

TRACE:
TRansport **AI** **C**arbon **E**valuator

As AI applications are rapidly rolled out at scale in all areas of the transport sector, decision-makers have no consistent, robust or timely means to estimate carbon impacts of AI and consider them alongside the applications' intended benefits.

AI in transport

Understanding the opportunities and environmental trade-offs

AI is already being applied across a wide range of transport activities, both in the transport sector and within DfT's own organisational functions.

Across the sector, AI is already being applied in traffic management, public transport operations, freight and logistics and infrastructure monitoring, where it is being used to improve real-time decision-making, optimise services, anticipate and manage risks and support more proactive management of assets and networks.

Within DfT and related organisations, there is a growing use of AI in internal decision support and administration, including the analysis of large volumes of text, automation of image-based tasks, faster handling of correspondence and more efficient review of consultation responses.

AI could play an important role in reducing carbon emissions across the economy, including in transport. The literature review conducted as part of this commission [1] highlights evidence that AI may support lower-emission outcomes through areas such as transport efficiency gains, improved forecasting, smarter system optimisation, better energy-grid management and more efficient industrial and logistics operations.

At the same time, **AI also has an environmental footprint that is significant and growing**. There are many concerns including high energy consumption, increasing data-centre emissions, the embodied carbon of hardware and infrastructure, water use for cooling, resource depletion linked to critical minerals and increased e-waste. These impacts are especially pronounced for more compute-intensive forms of AI, particularly generative AI and large language models.

While certain studies indicate AI's potential emissions savings can significantly outweigh its own emissions, the effectiveness of this can be undermined by **rebound effects**. This is when efficiency improvements save carbon in the short-term but make it easier to do more of an activity potentially negating the carbon savings generated. Such rebound effects are difficult to quantify but are a key determinant of the final outcome.

The implication for DfT is that AI should be approached neither as inherently beneficial nor inherently harmful. Its environmental value depends on how and where it is used. That creates a clear need for a practical way to assess likely benefits and carbon costs on a case-by-case basis.

Lack of visibility of carbon in decision-making

Challenges in understanding carbon impacts at early stages

As AI applications are brought forward across transport, decision-makers need a way to understand not only what an application might achieve, but also the environmental cost of deploying it.

A major challenge is that AI's environmental impacts are unusually difficult to quantify. Key data on hardware manufacture, model training and inference are often **proprietary**, especially where commercial cloud providers are involved.

Even where data exists, it is difficult to know where training or inference is **physically taking place**, what electricity mix is being used or how emissions should be allocated across **shared data-centre infrastructure**.

Carbon impacts can also vary significantly depending on inference volume (ie level of usage of an AI application), **usage patterns** and model complexity.

There is also weak **standardisation in assessment methods**, especially for embodied emissions and prompt-level inference costs.

Gathering the necessary data and processing it for each use case would also take **significant time**, making it difficult to support decisions at the pace applications are being developed.

This means that producing a **single, highly accurate lifecycle carbon number for every proposed AI application is not practical** for DfT. Yet decisions still need to be made, often at pace and while applications are still at an early stage of development.

There therefore needs to be a way for teams to undertake a consistent, proportionate and transparent assessment of carbon impacts early enough to support decision-making.

This requires a focus on decision-useful accuracy rather than theoretical precision, ensuring estimates are robust enough to inform early decisions despite uncertainty.

I have no idea what the environmental impact of different AI tools are at the moment. I'm going in blind and if I'm not informed, I can't make informed decisions.

Head of AI, arm's length body

The solution:
Introducing TRACE

TRACE is a first-of-its kind tool enabling policy- and decision-makers to consider the carbon costs of AI applications ahead of deployment in transport.

TRACE:
TRansport **A**I
Carbon **E**valuator

Introducing TRACE

Supporting decision making on the carbon impacts of AI in transport

What is TRACE?

In response to these challenges, TRACE provides a practical way for policy-makers and decision-makers to assess the carbon implications of AI-enabled interventions in the transport sector. The tool supports a high-level, whole-life carbon assessment, allowing users to compare an AI-enabled use case against a realistic counterfactual and consider whether its environmental cost is justified by its intended benefits.

In doing so, three key questions are addressed:

- What is the difference in carbon impact between an AI use case and counterfactual?
- How does the carbon impact of one use case compare to another?
- Do the benefits from the AI use case justify its carbon impact?

The tool takes a system-wide approach, incorporating emissions across defined system boundaries. It considers factors such as:

- AI hardware (embodied carbon from manufacturing)
- AI software (energy flowing through hardware stock)
- Transport activity (eg fleet deployment)
- Energy in field equipment
- Embodied carbon in purchased materials

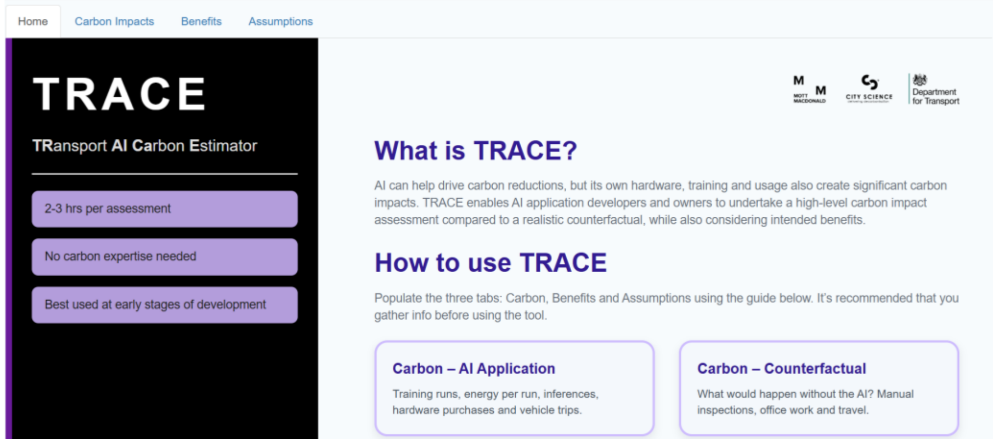

TRACE Operations

The tool is underpinned by a structured framework of emissions activities, which forms the core of its analytical approach. These activities represent the key sources of greenhouse gas (GHG) emissions associated with both AI-enabled systems and the conventional processes they may replace.

Methodology

All calculations are grounded in a standardised methodology aligned with the principles of the GHG Protocol, ISO 16404, ensuring that outputs are consistent with internationally recognised approaches to carbon accounting.

Core components of TRACE

How does TRACE support decisions on the application of AI in transport?

The tool provides a structured, system-wide comparison that enables decision-makers to assess whether the intervention is justified.

AI use case emissions

This section captures the emissions directly associated with the AI-enabled intervention itself, across its full lifecycle. It includes upstream infrastructure (embodied emissions from manufacturing compute hardware), training and fine-tuning (energy used to develop and adapt models), inference (ongoing operational energy use as the model generates outputs) and end-user terminals (any additional devices required to enable the use case).

All other emissions

This section captures the broader system impacts beyond core AI compute, ensuring that the assessment reflects the full real-world context of the intervention. It includes both the emissions associated with delivering the AI-enabled use case and those associated with the counterfactual scenario it replaces. Examples may include operating a vehicle with an AI sensor or purchasing construction materials for an infrastructure intervention.

Sensitivities

Sensitivity analysis allows users to test how results change under different assumptions for key parameters such as grid carbon intensity, manufacturing complexity, data centre lifetime and average utilisation. This helps decision-makers understand how dependent the results are on underlying assumptions and whether the environmental case for using AI remains robust under more or less favourable conditions.

Benefits

This section captures the intended outcomes of the AI-enabled intervention, enabling users to assess whether the environmental costs are justified by the value delivered. Benefits are defined by the user in relation to the functional objective of the use case and may include improvements such as reduced travel time, increased network efficiency, enhanced safety or operational cost savings.

Tool structure

Inputs and outputs

The tool draws on a set of user inputs and produces high-level outputs that assess whole-life carbon impacts. The inputs capture key assumptions about the AI use case and operational context, while the outputs translate these into carbon emission estimates.

The tool has six components (shown below), each with a clear purpose that contributes to the decision-making process.

1 Use case

To establish a clear definition of the use case and its current operating baseline.

2.A AI model

To establish the model type and complexity of AI system being used.

2.B Counterfactual

To identify the alternative method capable of achieving the same functional outcome as the proposed AI system.

3 Carbon impact

To establish an indicative carbon impact profile created by the intervention, covering energy usage and embodied emissions in both scenarios.

4 Sensitivity testing

To allow testing of the major assumptions underpinning the tool such as grid carbon intensity.

5 Benefits

To establish an indicative benefits profile created by the AI system relative to its carbon impact.

6 Outputs

To integrate carbon impact and benefit results to support decision-making.

How accurate are the results? Should you trust them?

Interpreting results for early decision-making

The tool produces **robust, policy-relevant estimates** by combining user-defined activity data with a structured emissions factor database aligned to established carbon accounting principles, enabling consistent whole-system carbon comparisons across AI use cases.

However, like all lifecycle carbon assessments, results reflect **inherent uncertainty in complex systems such as AI**, where supply chains, compute locations, and usage patterns are not fully observable or standardised in practice.

In line with the GHG Protocol principles of completeness, transparency, and practicality, the tool prioritises **decision-useful accuracy over theoretical precision**, recognising that perfect measurement is neither feasible nor necessary for robust policy appraisal.

This approach reflects established carbon accounting practice that models and estimates are valid and decision-grade when transparently constructed, even where they rely on well-justified assumptions rather than full primary measurement.

The outputs should therefore be interpreted as **best-available, structured estimates** that enable consistent comparison of options, rather than exact measurements of emissions.

The value of the tool is not in achieving perfect precision, but in taking the first structured step towards quantifying the carbon impacts of AI in transport, enabling robust, informed decision-making today while creating a foundation for progressively greater accuracy and sophistication in future iterations.

This tool is therefore a **vital first step** in bringing carbon considerations into AI-related transport decision-making, acknowledging that – while initial estimates may not have perfect precision – they will still be decision-relevant.

This structured baseline will **create the foundation** for progressively improving accuracy, consistency, and confidence in future iterations as data and methods mature.

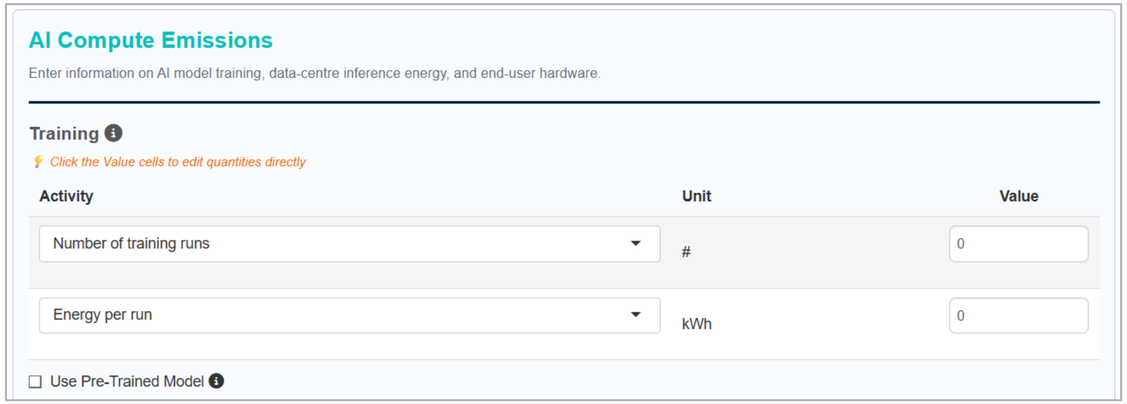

Using TRACE
Who should use it and how?

This section details a quick guide outlining the key inputs required to use TRACE and the resulting outputs, which are designed to support interpretation and inform decision-making.

TRACE:
TRansport **AI**
Carbon **E**valuator

Quick start guide

How to use TRACE

For first-time users, the tool offers a simple way to assess carbon impacts at an early stage. It is designed to be completed in two to three hours, jointly by colleagues in policy, focusing on the problem and AI tool is trying to solve and data science, as well as the technical aspects of the AI application to be deployed. To generate an emissions profile, users provide information (as shown below) describing the scale of relevant activities within their project.

What is required?

Policy Inputs | **Data Science Inputs**

Core inputs: These inputs establish the context within which emissions activities are selected and quantified.

Use case overview: Users define the scope of their AI application and the corresponding counterfactual scenario, establishing the basis for comparison.

AI compute and other use case emissions: Users enter activity data across the AI lifecycle (eg compute, deployment) as well as other relevant emissions-generating activities within the defined system boundary.

Counterfactual emissions: Users input emissions associated with the counterfactual scenario, representing what would occur without the AI use case (eg vehicle operation).

Sensitivity: Users set parameters based on chosen assumptions, referring to the benchmarking visuals to inform inputs.

Benefits: Users can assess the relative advantages of the AI use case compared to the counterfactual by providing inputs on impact, confidence and importance.

The images below illustrate examples of the required inputs for carbon impacts:

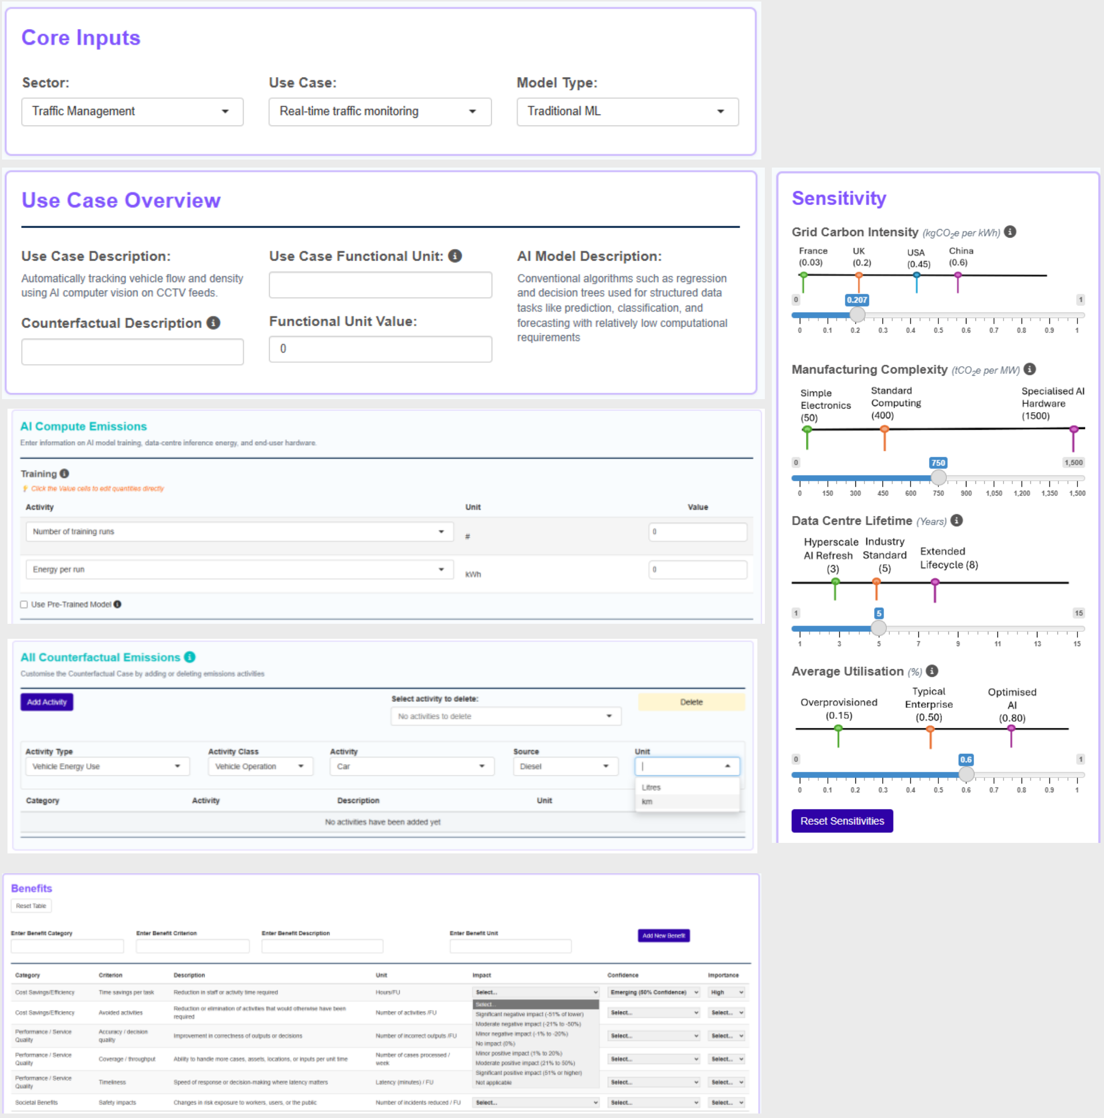

Case Study

Tool outputs and how they inform decision-making

Initial tool outputs: The initial outputs from the tool include a structured results table comparing carbon emissions between the AI-enabled scenario and the counterfactual. These results are derived from the key assumptions underpinning the assessment [2], to ensure transparency and traceability (eg Embodied Carbon of ICT Infrastructure).

In addition, the tool generates a benefits visualisation in the form of a graph, which evaluates performance across each assessment criterion. This provides a clear, criterion-by-criterion view of impacts, indicating whether the use of AI is beneficial in each area. Rather than aggregating results into a single composite metric, this approach preserves visibility of individual impacts, highlighting both areas of benefit and potential disbenefit and supporting a more nuanced interpretation of outcomes.

Downloadable tool outputs: TRACE also enables users to export results into a single consolidated document designed to support decision-making processes and governance forums (illustrative example shown in Appendix A). The output provides a clear summary of the assessed carbon impacts and benefits, alongside key assumptions and contextual information to support interpretation of the results.

How does this feed into decision-making? The outputs are intended as a "tool for thinking" rather than a definitive answer. They provide structured, comparable information on carbon impacts and underlying assumptions, helping users explore trade-offs, test scenarios and support more transparent and evidence-based decision-making on AI deployment.

The images below illustrate examples of the tool outputs:

Carbon Impacts Results:

Generate Results

Download Results (Excel) Download Results (Word)

Case	Category	Activity	Use Case (tCO_2e)	Counterfactual (tCO_2e)
AI Usage	Traditional ML	Data Centre & ICT Equipment Manufacture	1.41	0.00
AI Usage	Traditional ML	Training	6.21	0.00
AI Usage	Traditional ML	Inference	3.10	0.00
AI Usage	Traditional ML	End User Terminal Manufacture	0.01	0.00
Other Emissions	Purchasing of Equipment	Laptop Purchase	0.01	0.00
Other Emissions	Vehicle Operation	Car	0.00	12.85
Total			10.75	12.85

Benefits Results:

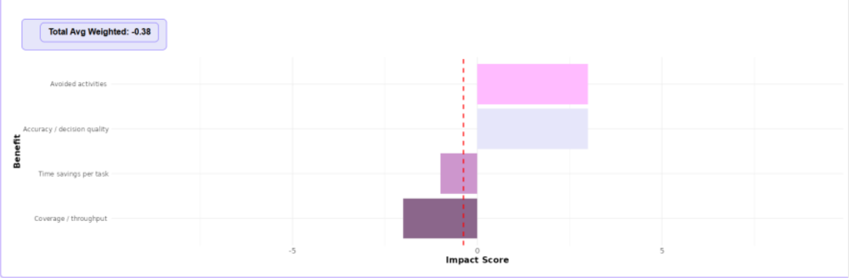

Downloadable Excel Tool Output:

	A	B	C	D	E
1	Case	Category	Activity	Use Case (tCO_2e)	Counterfactual (tCO_2e)
2	AI Usage	Traditional ML	Data Centre & ICT Equipment Manufacture	1.41	0.00
3	AI Usage	Traditional ML	Training	6.21	0.00
4	AI Usage	Traditional ML	Inference	3.10	0.00
5	AI Usage	Traditional ML	End User Terminal Manufacture	0.01	0.00
6	Other Emissions	Purchasing of Equipment	Laptop Purchase	0.01	0.00
7	Other Emissions	Vehicle Operation	Car	0.00	12.85
8	Total			10.75	12.85
9					

Immediate impact:

Benefits to DfT and the transport sector

TRACE:
TRansport **AI**
Carbon **E**valuator

DfT is considering TRACE forming part of its AI decision-making processes, supporting its statutory requirement to consider the environment and providing an opportunity to drive significant carbon reduction.

Immediate impacts

TRACE supports stronger governance of AI projects by producing clear, decision-useful information on their environmental impacts.

The tool generates a report that helps decision-makers understand the carbon implications of an AI application alongside its value, delivery considerations and risks. This supports DfT's statutory requirement to consider the environment when deploying AI and helps ensure decisions take account of whether the expected benefits of an AI tool justify the carbon used, particularly for significant AI projects where carbon use or energy demand may be material.

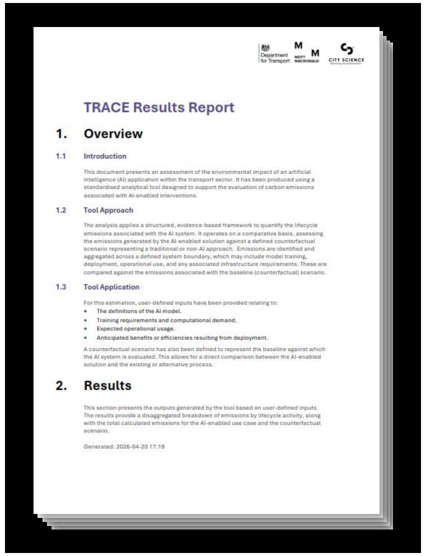
Department for Transport MOTT MACDONALD CITY SCIENCE

TRACE Results Report

1. Overview

1.1 Introduction

This document presents an assessment of the environmental impact of an artificial intelligence (AI) application within the transport sector. It has been produced using a standardised analytical tool designed to support the evaluation of carbon emissions associated with AI-enabled interventions.

1.2 Tool Approach

The analysis applies a structured, evidence-based framework to quantify the lifecycle emissions associated with the AI system. It operates on a comparative basis, assessing the emissions generated by the AI-enabled solution against a defined counterfactual scenario representing a traditional or non-AI approach. Emissions are identified and aggregated across a defined system boundary, which may include model training, deployment, operational use, and any associated infrastructure requirements. These are compared against the emissions associated with the baseline (counterfactual) scenario.

1.3 Tool Application

For this estimation, user-defined inputs have been provided relating to:

- The definitions of the AI model.
- Training requirements and computational demand.
- Expected operational usage.
- Anticipated benefits or efficiencies resulting from deployment.

A counterfactual scenario has also been defined to represent the baseline against which the AI system is evaluated. This allows for a direct comparison between the AI-enabled solution and the existing or alternative process.

2. Results

This section presents the outputs generated by the tool based on user-defined inputs. The results provide a disaggregated breakdown of emissions by lifecycle activity, along with the total calculated emissions for the AI-enabled use case and the counterfactual scenario.

Generated: 2026-04-20 17:19

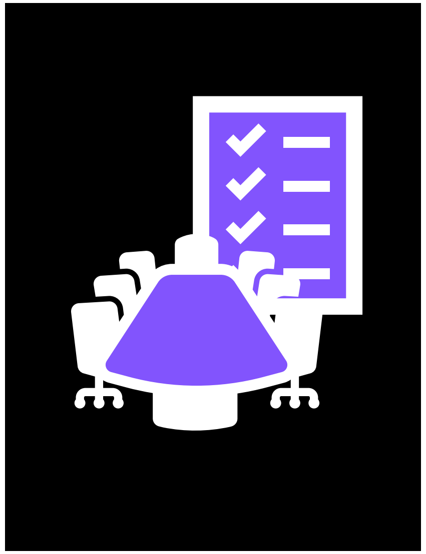

Carbon impacts can now be consistently included in decision-making for AI tools, alongside value

Output

Report generated by TRACE.

Outcome

Decision-useful information on carbon provided as standard to the **DfT AI Project Oversight Group**.

Impact

Increased visibility of carbon in decision-making when AI applications are rolled out in DfT will ensure better value is achieved for carbon used.

Initial lessons learnt

Lessons and trends arising from development and use of TRACE

Comparative analysis: AI-enabled interventions are not inherently more carbon intensive than their counterfactuals; rather, outcomes are highly contingent on system-level effects, where efficiency gains (eg reduced transport activity, optimised resource use) can outweigh the additional emissions associated with AI compute. This reflects a broader finding in the literature that digital optimisation can lead to net emissions reductions where it displaces traditional processes at scale, even when the enabling technology itself is energy-intensive.

Model family: Model architecture and complexity are primary drivers of variability in AI emissions, with more computationally intensive approaches – particularly generative AI and large-scale transformer models – exhibiting substantially higher energy requirements (often in orders of magnitudes in the millions) across training, fine-tuning and inference. These differences propagate upstream, as more complex models require larger volumes of specialised hardware, which also increases embodied emissions associated with manufacturing advanced ICT equipment.

System-level trade-offs: Model selection creates system-level trade-offs between carbon intensity and functional capability. particularly when comparing more computationally intensive approaches such as generative AI with less intensive methods. More capable models may increase AI-related emissions but also imply a more carbon-intensive counterfactual, as achieving the same functional outcome through conventional means would require greater activity. Conversely, simpler models reduce AI emissions but may deliver smaller system-level savings. The optimal solution depends on whether increased AI emissions are offset by reductions elsewhere in the system.

Procurement: Emissions associated with purchased goods and infrastructure can, in many cases, dominate total lifecycle impacts within the system boundary, particularly where interventions involve significant material inputs (eg vehicles, hardware or physical assets). In line with lifecycle assessment and GHG Protocol guidance, these cradle-to-gate emissions are fully attributable within the assessment boundary, meaning that variations in material demand between the use case and counter-factual can be a more significant driver of overall emissions than AI-related energy use alone.

Appendix A: Sample outputs from TRACE

When TRACE is completed, it can generate a Word document containing a full record of the carbon, benefits and assumptions.

TRACE:
TRansport **AI**
Carbon **E**valuator

Sample outputs from TRACE

TRACE Results Report

1. Overview

1.1 Introduction

This document presents an assessment of the environmental impact of an artificial intelligence (AI) application within the transport sector. It has been produced using a standardised analytical tool designed to support the evaluation of carbon emissions associated with AI-enabled interventions.

1.2 Tool Approach

The analysis applies a structured, evidence-based framework to quantify the lifecycle emissions associated with the AI system. It operates on a comparative basis, assessing the emissions generated by the AI-enabled solution against a defined counterfactual scenario representing a traditional or non-AI approach. Emissions are identified and aggregated across a defined system boundary, which may include model training, deployment, operational use, and any associated infrastructure requirements. These are compared against the emissions associated with the baseline (counterfactual) scenario.

1.3 Tool Application

For this estimation, user-defined inputs have been provided relating to:

- The definitions of the AI model.
- Training requirements and computational demand.
- Expected operational usage.
- Anticipated benefits or efficiencies resulting from deployment.

A counterfactual scenario has also been defined to represent the baseline against which the AI system is evaluated. This allows for a direct comparison between the AI-enabled solution and the existing or alternative process.

2. Results

This section presents the outputs generated by the tool based on user-defined inputs. The results provide a disaggregated breakdown of emissions by lifecycle activity, along with the total calculated emissions for the AI-enabled use case and the counterfactual scenario.

Generated: 2026-04-20 17:19

Case	Category	Activity	Use Case (tCO_2e)	Counterfactual (tCO_2e)
AI Usage	Traditional ML	Data Centre & ICT Equipment Manufacture	0.19	0.00
AI Usage	Traditional ML	Training	0.21	0.00
AI Usage	Traditional ML	Inference	1.03	0.00
AI Usage	Traditional ML	End User Terminal Manufacture	0.51	0.00
Other Emissions	Vehicle Operation	Car	2.07	5.19
Total			**4.01**	**5.19**

3. Selected Inputs

3.1 Core Inputs

The following core inputs have been selected:

- Sector: Traffic Management
- Use Case: Real-time traffic monitoring
- AI model Type: Traditional ML

3.2 Use Case Overview

The entered definitions of the chosen AI use case are set out below:

- The use case is described as: Automatically tracking vehicle flow and density using AI computer vision on CCTV feeds.
- The counterfactual is described as: Manual inspection
- The functional unit is described as: km of road inspected
- The functional unit value entered is 5000

3.3 Use Case Emissions

3.3.1 AI Compute Emissions

These AI compute emissions arise from the energy required to train and operate AI models, as well as the manufacturing of dedicated hardware required to deploy them. Corresponding to each stage of the AI lifecycle the following inputs have been entered for training:

*For this particular use case, the pre-trained model option has been *enabled.

Activity	Unit	Value
Number of training runs	#	0
Energy per run	kWh	0

Once deployed, AI models generate outputs by performing inference, which refers to the process of applying a trained model to new input data. Inference workloads typically occur at a much higher frequency than training events and can therefore represent a significant share of operational emissions over the lifetime of an AI system. The following inference inputs for the chosen use case have been entered:

The following inference inputs for the chosen use case have been entered:

Activity	Unit	Value
Number of inferences	#	1,000,000.00 0
Energy per inference	kWh	0.005

3.3.2 Fine-Tuning

Training an AI model involves running computational workloads across specialised hardware, typically located in cloud or data-centre environments. The energy required for this process can be substantial depending on the size and complexity of the model.

The following fine-tuning inputs for the chosen use case have been entered:

Activity	Unit	Value
Total energy consumption per fine-tuning event 1	kWh	1,000

3.3.3 End-User Terminal

End-user terminals inlcude only new or additional hardware procured specifically to enable the AI use case. The selected end-user terminals are shown below:.

Activity	Unit	Value
Camera	Number	32
Edge AI Module	Number	4

3.3.4 Other Use Case Emissions

Other use case emissions inputs are shown in the table below.

Sample outputs from TRACE – continued

Activity	Description	Source	Unit	Value
Vehicle Operation – Car	2x Inspected distance	Electricity	km	10,000

3.4 Counterfactual

Counterfactual activities are shown in the table below.

Activity	Description	Source	Unit	Value
Vehicle Operation – Car	6x inspected distance	Diesel	km	30,000

4. Benefits

Benefits are important to this assessment because it helps to assess the advantages of using an AI-enabled use case and a corresponding counterfactual scenario. The primary purpose is to support users in understanding the value created by the AI system through structured comparison across key criteria. The output provided is intended to support comparative decision making rather than to produce definite benefit estimates.

- Impact: The extent to which this outcome is expected to improve or worsen compared to a scenario without using AI.
- Confidence: The level of certainty associated with this estimate.
- Importance: The relative significance of this benefit within the overall assessment.

The benefits inputs provided in the tool are shown below.

Category	Criterion	Description	Unit	Impact	Confidence	Importance
Cost Savings/Efficiency	Time savings per task	Reduction in staff or activity time required	Hours/FU	Significant positive impact (51% or higher)	Proven (100% Confidence)	High
Cost Savings/Efficiency	Avoided activities	Reduction or elimination of activities that would otherwise have been required	Number of activities /FU	Significant positive impact (51% or higher)	Proven (100% Confidence)	High
Performance /	Accuracy / decision quality	Improvement in correctness of outputs or decisions	Number of incorre	Moderate positive impact	Emerging (50% Confidence)	Moderate

Department for Transport | Mott MacDonald | City Science

Criterion	Impact	Impact.Score	Confidence	Confidence.Score	Importance	Importance.Score	Weighted.Score
	(21% to 50%)						
Coverage / throughput	Significant positive impact (51% or higher)	3	Established (75% Confidence)	0.75	Moderate	2	4.5
Timeliness	Moderate positive impact (21% to 50%)	2	Established (75% Confidence)	0.75	Low	1	1.5
Safety impacts	Moderate positive impact (21% to 50%)	2	Established (75% Confidence)	0.75	Moderate	2	3.0

5. Assumptions

5.1 AI Compute Emissions Assumptions

The following assumptions have been entered for AI compute emissions:.

num	Category	Input	Unit	Assumption
1	Training	Number of training runs	#	Pre-trained model
2	Training	Energy per run	kWh	Pre-trained model
3	Inference	Number of inferences	#	2,000 inferences per km
4	Inference	Energy per inference	kWh	0.005 based on [source]
5	Fine-Tuning	Total energy consumption per fine-tuning event 1	kWh	1000 based on [source]
6	End-User Terminal	Camera	Number	8 cameras per vehicle for 4 vehicles
7	End-User Terminal	Edge AI Module	Number	1 edge device per

Appendix B: Development journey

Research to support this decision-useful tool

TRACE:
TRansport **A**I
Carbon **E**valuator

This section sets out the research, evidence and analytical process that informed the development of TRACE. It outlines the key insights, challenges and methodological choices that shaped the tool, drawing on technical material developed across the commission.

What led to the development of TRACE

TRACE arose from a DfT Futures project [3] to explore the carbon impacts of AI

The Problem

AI has real promise to enable better transport for all, but this comes with trade-offs, which are not currently visible in AI decision-making.

The rapid expansion of generative AI and large-scale model training is driving up energy demand, which risks undermining potential gains.

However, we don't currently understand the scale of carbon impacts and what factors influence emissions.

Futures Study: Carbon Impacts of AI

A Futures study used a literature review and seven-questions interviews to gather evidence on AI's lifecycle carbon emissions and the needs of teams working with AI.

This work gathered evidence on the lifecycle carbon costs of machine learning across embodied, training and inference stages, identified a wide range of AI use cases for further testing, and highlighted key areas of uncertainty in the future carbon impacts of AI.

The research highlighted the need for an 'at-pace' approach. Recognising uncertainty, data variability and a fast-changing AI landscape, informed decisions were prioritised over highly precise carbon calculations.

Understanding these needs alongside the complexity of AI lifecycle emissions led to the proposal to develop a decision-making tool to support policy- and decision-makers.

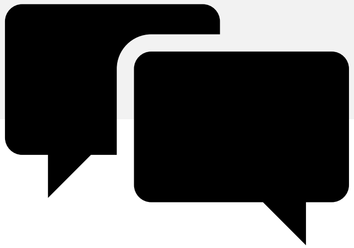

Lifecycle Analysis: Overview

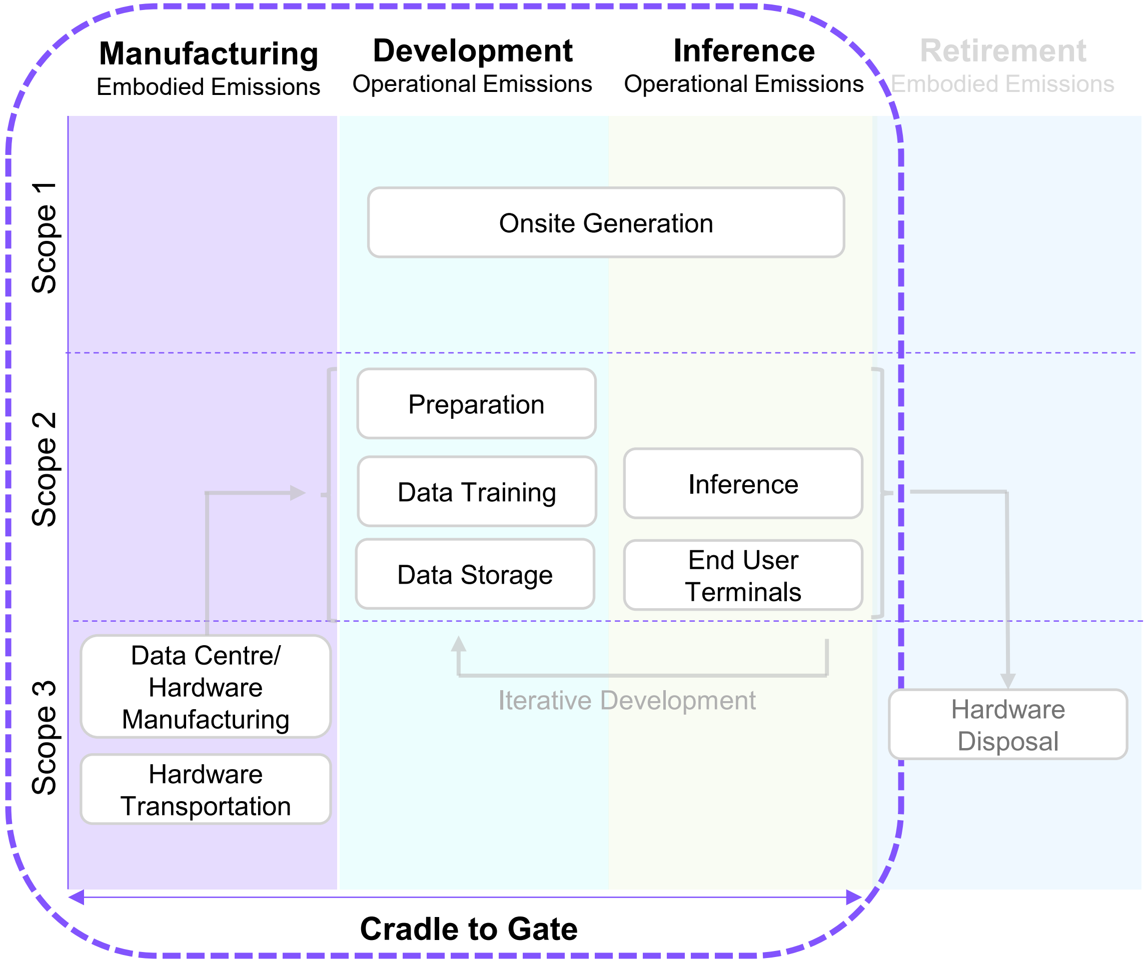

Based on literature review

The Futures study that preceded TRACE used Life Cycle Analysis (LCA) to map and quantify the Greenhouse Gas (GHG) emissions from AI models. Our boundary is **cradle-to-gate**: starting from resource extraction ("cradle") and concluding when the AI service is ready for use ("gate").

The three core phases of the AI lifecycle are: Manufacturing (Embodied Carbon), Use (Operational Carbon), and Retirement (End-of-Life). We conduct this analysis by breaking these lifecycle stages down into fundamental LCA building blocks:

- **Unit Process**: The distinct, physical building blocks of the system. In the chart, these are the hardware phases, *e.g. "Data Centre/Hardware Manufacturing," "Onsite Generation," and "End User Terminals"*.
- **Economic Flow**: The material inputs (goods, services, energy) that flow between Unit Processes and drive the consumption of resources. These are the software activities in the chart, *e.g. "Data Training," "Inference," and "Data Storage"*
- **Functional Unit:** The reference by which all inputs and outputs are normalised, or the intensity variable. *E.g. "per inference (or query)"*

The model is constructed by quantifying these economic flows between the unit processes that make up the system boundary, in this case, the cradle-to-gate lifecycle of AI models used in transport applications.

The goal is not to use LCA to produce perfect numbers, but consistent, comparative ones that support decisions.

Note: this approach excludes the end-of-life (cradle-to-grave) phase. This system boundary is strategically aligned with the GHG Protocol's Scope 3 guidance, specifically Category 1: Purchased Goods and Services.

Lifecycle analysis: mapping the basics

The lifecycle of AI is less intuitive than other traditional industries because it includes both **hardware** and **software** phases that interact dynamically. To calculate environmental impacts, we must precisely map these processes to one another, ensuring all activity within the system boundary is identified and quantified. Conceptually, the hardware lifecycle represents our unit processes, while the software lifecycle defines the necessary economic flows operating *through* those units. Each software stage is mapped and allocated to its corresponding hardware and cumulatively forming the basis of the cradle-to-gate emissions. To build a meaningful LCA, we:

- **Defined the hardware lifecycle:** the complete physical infrastructure necessary to support the AI model, encompassing manufacture (embodied) and operation (use).
- **Defined the software lifecycle:** the activities that drive demand on hardware and consume energy, and thus emissions.
- **Mapped the two together:** systematically linking software (economic flows) to hardware (unit processes) across all relevant phases (e.g. manufacture, training, inference).

Hardware lifecycle
(unit processes)

- Data centre – computer hardware (GPU/CPU)
- Data centre – memory/storage (SSD, DRAM)
- Data centre – networking (switches, routers)
- Data centre – cooling & infrastructure (auxiliary)
- Web hosting
- Networks
- End-user terminals

Software lifecycle
(economic flows)

- Preparation: The initial stage of gathering data and selecting algorithms to set up an AI model for training.
- Data Training: The process of feeding data into an AI model to adjust its parameters and improve performance.
- Data Inference: Using a trained AI model to make predictions or decisions on new, unseen data.
- Data Storage: The act of saving datasets, model parameters, and outputs for future access and use.

Lifecycle Phase

1. Manufacturing*
2. Training
3. Interference
4. End of Life

**Including transport*

Allocation

Allocation is a key technical challenge in calculating the lifecycle environmental impact of AI. Economic flows and unit processes do not always align perfectly:

- Shared servers and cooling systems make it difficult to determine the volume of energy relevant to a single model's economic flow.
- Many data centre processes produce outputs outside the defined system boundary

A credible and actionable calculation will therefore require simplifying assumptions and proxies to inform decision-making.

Lifecycle analysis: full boundary mapping

System location	Hardware (unit process)	Software/activity (economic flow)	Classification
Data centre	Cooling & infrastructure	Manufacturing	*Embodied: long-term lifecycle*
		Cooling	*Use: short-term marginal*
	Memory & storage (SSD, DRAM)	Training	*Embodied: short-term marginal*
		Manufacturing	*Use: short-term marginal*
		Inference	*Use: short-term marginal*
		Storage	*Embodied: long-term lifecycle*
	Computer hardware (GPU/CPU)	Manufacturing	*Use: long-term lifecycle*
		Training	*Use: long-term lifecycle*
		Inference	*Use: short-term marginal*
	Networking (switches, routers)	Manufacturing	*Embodied: long-term lifecycle*
		Training	*Use: long-term lifecycle*
		Inference	*Use: short-term marginal*
Web hosting	Web hosting (cloud)	Manufacturing	*Embodied: long-term lifecycle*
		Inference	*Use: short-term marginal*
Networks	Network infrastructure	Manufacturing	*Embodied: long-term lifecycle*
		Inference	*Use: short-term marginal*
End-user terminals	User platform	Manufacturing	*Embodied: long-term lifecycle*
		Inference	*Use: short-term marginal*

Presented here is a comprehensive mapping developed by our team over the course of the study. This lifecycle synthesised from multiple studies of the hardware, software, and lifecycle stages used to model a robust, cradle-to-gate emissions impact of AI applications in transport.

This structured mapping identifies where each process sits within the lifecycle, its measurement boundaries, and its inclusion criteria. While highly granular studies exist, such as those separating training into dozens of sub-phases, that level of detail is not feasible for scalable, practical measurement. Conversely, an overly aggregated model would lose decision-useful accuracy. We believe this framework strikes a necessary balance: detailed enough to be representative, yet feasible to populate with available data.

However, calculating a definitive, single lifecycle number for every case based on this extensive detail may not feasible. It is difficult to ascertain a single lifecycle number because of data access limitations, rapid technological development, the impossibility of tracking precise geographical location for every inference, and the technical challenge of allocating shared energy usage within data centre hardware.