




► **Project Symbiosis**
Part 1: Key Findings

Exploring AI for scope 3 accounting and transition finance

October 2025

Contents

Abstract	3
1 Introduction	4
2 Scope 3 emissions	5
Accounting standards	5
Calculation methods	6
Calculation challenges	7
Knock on implications	8
Impetus to explore the use of AI and related technologies	9
3 Project Symbiosis: exploring AI and related technologies for scope 3 accounting and transition finance	11
Overview	11
User groups	12
Explored approaches and outcomes	13
Data collection	13
Calculation and reporting	14
Reduction analysis	14
Applied example	16
4 Learnings and next steps	17
Acknowledgements	19

Abstract

Project Symbiosis – a collaboration between the BIS Innovation Hub Hong Kong Centre and the Hong Kong Monetary Authority (HKMA) – performed applied technology research, including through developing the Net Emissions Optimiser (NEMO) proof-of-concept (POC), to showcase how novel technologies offer a viable technical pathway to positively impact core stakeholders, including corporations, the financial sector, small and medium-sized enterprises (SMEs), people and the planet by reducing critical information gaps impeding the climate transition. The research was set against the backdrop of challenges unique to the financial sector, where over 95% of emissions fall in the scope 3 category, and the HKMA's Sustainable Finance Action Agenda.

As demonstrated through the work of Project Symbiosis, AI and related technology approaches have the potential to improve the status quo, whether by improving the speed, breadth and quality of data collection, by generating intelligent environmental impact results that are flexible with respect to a wide range of data availability scenarios, or by generating financeable emission reduction opportunities. Collectively, the explored approaches demonstrate real potential for AI to help bridge knowledge gaps that are currently impeding the provision of needed transition finance.

The learnings of Project Symbiosis, while applied in the context of the use case of scope 3 reporting, are of broader significance, illustrating how AI and related technologies may be applied in a range of contexts such as risk assessment, supply chain management, and reporting automation as its rapid evolution continues.

The project name Symbiosis was chosen because more accurate calculation of scope 3 emissions and impact data, combined with funding sources to reduce them, could achieve a symbiotic relationship between core stakeholders in complex supply chains.

The image of the clownfish was chosen as it lives in a symbiotic bond with the sea anemone that it inhabits, and the transition finance matching engine was likewise named the Novel Emissions Optimiser or NEMO.

While poisonous to other fish, the anemone offers protection to the clownfish that live in it.¹

1 See [Friends and Anemones: How Clownfish Strengthen Symbiotic Bonds with Their Hosts](#) | Discover Magazine 2025.

I. Introduction

Many companies do not directly produce the goods they offer to consumers. Rather, they acquire those goods under contract from third-party suppliers. This creates a challenge to understanding not only the greenhouse gas (GHG) emissions profile of the goods, but also the actions that can be taken to reduce emissions and impacts. Companies struggle to obtain sufficient data to calculate the emissions that are “upstream” and “downstream” of their direct business operations (ie their scope 3 emissions), which often make up more than 90% of their overall company emissions.

On the other end of the value chain are small and medium-sized enterprises (SMEs), which often operate several tiers (and continents) away from the companies required to undertake emissions accounting, and are even further removed from the financial institutions that can provide transition finance. Although critical emissions reduction opportunities exist for these smaller and removed enterprises (eg product manufacturers, material producers), their emissions are often not directly addressed. The result may be limited awareness or incentive to undertake retrofits to equipment that otherwise function (eg fossil fueled or energy inefficient equipment or power sources).

Breakthrough technology solutions are needed that can bridge these gaps while at the same time leveraging the untapped power of financial institutions to accelerate the climate transition.

Against this backdrop, Project Symbiosis - a collaboration between the BIS Innovation Hub Hong Kong Centre and the Hong Kong Monetary Authority – set out to assess whether AI and related technology approaches can be part of the solution.

- The first goal of Project Symbiosis is to explore how advanced data techniques and AI can be leveraged to more accurately collect, interpret and calculate scope 3 emissions and other impact data in corporate supply chains.
- The second goal of Project Symbiosis is to explore how AI can identify emissions reduction opportunities based on such data and thereby lay the groundwork for reducing scope 3 emissions.
- Finally, Project Symbiosis aims to explore how to leverage emissions and reductions data to match suppliers with funding sources to decarbonise the supply chain (referred to as financeable emissions reduction opportunities).

Project Symbiosis’ applied technology research developed amongst others the Net Emissions Optimiser (NEMO) proof-of-concept (POC). The research was set against the backdrop of challenges unique to the financial sector, where over 95% of emissions fall in the scope 3 category, and the HKMA’s Sustainable Finance Action Agenda.²

² Hong Kong Green Finance Association, HKMA unveils Sustainable Finance Action Agenda, October 2024, <https://www.hkgreenfinance.org/hkma-unveils-sustainable-finance-action-agenda/>.

The project demonstrates that AI and related technology approaches have the potential to improve the status quo, whether by improving the speed, breadth and quality of data collection, by generating intelligent environmental impact results that are flexible with respect to a wide range of data availability scenarios, or by generating financeable emission reduction opportunities. Collectively, the explored approaches demonstrate real potential for AI to help bridge knowledge gaps that are currently impeding the provision of needed transition finance.³

The work of Project Symbiosis builds upon the work completed in Project Gaia, which deployed novel AI-based solutions to assess climate-related financial risks based on publicly reported corporate environmental, social and governance (ESG) reports. Project Symbiosis focuses on the underlying sources of these data, exploring pathways for developing reportable emissions results and identifying emissions reduction opportunities in sectors and smaller scale businesses that are otherwise left out of the emissions reports produced by large multinational corporations. The recently launched project Danu builds on projects Gaia and Symbiosis by investigating the use of digital twin technologies, which in turn leverage AI and related technologies to measure physical risks and assess their impact on financial stability.

The learnings of Project Symbiosis, while applied in the context of the use case of scope 3 reporting, are of broader significance, illustrating how AI and related technologies may be applied in a range of contexts such as risk assessment, supply chain management, and reporting automation as its rapid evolution continues.

II. Scope 3 emissions

A. Accounting standards

Scope 3 emissions refer to the indirect greenhouse gas emissions that occur in the value chain of an organisation from upstream and downstream activities not under their ownership or control. This framework of emissions scopes derives from the **Greenhouse Gas (GHG) Protocol**, the leading standard for the calculation of organisation-level emissions.⁴

The GHG Protocol is embedded within major regulatory reporting requirements across the globe, including the **International Sustainability Standards Board (ISSB)**⁵

3 At the onset of the project, the BIS Innovation Hub convened a workshop exploring challenges and goals. Participants from over 20 organisations attended, including corporations, financial institutions, regulators and ecosystem enablers. Acknowledgements can be found at the conclusion of Part 1.

4 World Business Council for Sustainable Development (WBCSD) and World Resources Institute (WRI), The Greenhouse Gas Protocol, A corporate accounting and reporting standard, 2004, <https://ghgprotocol.org/corporate-standard>.

5 The ISSB was established by the International Financial Reporting Standards (IFRS) Foundation in 2021 to create a global baseline for sustainability-related financial disclosures. The ISSB aims to: standardize ESG reporting across jurisdictions and industries; help companies disclose decision-useful, investor-focused sustainability information; and consolidate existing frameworks like TCFD, SASB, CDSB, and the Integrated Reporting Framework. See IFRS - [International Sustainability Standards Board](https://www.ifrs.org/news-and-events/news/2024/05/jurisdictions-representing-over-half-the-global-economy-by-gdp-take-steps-towards-issb-standards/).


6 International Financial Reporting Standards (IFRS), Jurisdictions representing over half the global economy by GDP take steps towards ISSB standards, May 2024, www.ifrs.org/news-and-events/news/2024/05/jurisdictions-representing-over-half-the-global-economy-by-gdp-take-steps-towards-issb-standards/.

standards that are slated for adoption in jurisdictions representing more than half of the global economy,⁶ the EU Corporate Sustainability Reporting Directive (CSRD), and California's Senate Bills 253 and 261. These reporting expectations are forecast to grow over time, despite pressure to simplify reporting requirements in some reporting jurisdictions.

Across the global economy, scope 3 emissions represent roughly 75% of the emissions reported by companies. And for financial institutions, this number represents more than 95% of reported emissions, virtually all of which is concentrated in their scope 3, category 15 financed emissions. Thus, scope 3 accounting is of critical importance in addressing climate impacts, but challenges abound.

Graph 1.1

Overview of Scopes 1, 2 and 3 emissions



B. Calculation methods

The **GHG Protocol scope 3 standard** breaks down emissions that occur upstream and downstream of the company's directly owned and controlled operations into 15 distinct categories.⁷ While the scope 3 standard defines rigid boundaries for each of these categories, it permits significant flexibility in terms of calculation method, often suggesting more than three optional calculation approaches for each category, as well as the possibility of deploying hybrid methods.

For example, with respect to scope 3, category 1 (purchased goods and services), the GHG Protocol scope 3 standard and associated technical guidance permit a series of methods for quantifying emissions, including the:

- supplier-specific method (direct data from suppliers);

⁷ WBCSD and WRI, Corporate value chain (Scope 3) accounting and reporting standard, 2011, <https://ghgprotocol.org/corporate-value-chain-scope-3-standard>.

- hybrid method (combination of supplier-specific data and secondary data);
- average-data method (general data like product weights and secondary data); and
- spend-based method (multiplying total spend on qualifying products and services by spend-based secondary emissions factors).⁸

Calculating results using the same underlying package of information can result in significant differences in emissions results – both across the different approaches and within a given method, as companies often use different secondary data with different emissions estimates for the same materials and products.

The ultimate result of this flexibility is that corporate inventory data lack real comparability and in many cases produce very little visibility into actual activities causing scope 3 emissions.

C. Calculation challenges

Companies have widely divergent systems for managing supply chain data, almost always in systems not primarily designed for tracking and organising the key information required for sustainability reporting. The result is vastly inconsistent primary data availability across companies and even within corporate groups, in terms of quality, breadth and the type of data points used for calculation. These variances in input data further inhibit standardisation in calculation approaches.

In many cases, sustainability-focused functions within companies spend inordinate amounts of time collecting and organising such data for the purpose of emissions accounting and reporting, often with unresolvable gaps in their own internal primary data sets. As a result, businesses are often forced to rely on either inaccurate proxy data (eg the global average emissions for a raw material, rather than specific supply chain data for the actual raw material used) or, more commonly, spend-based accounting for scope 3.

Why spend-based accounting is imperfect

The most common fallback method for scope 3 emissions accounting is to multiply corporate spending by economic emissions factors (eg “\$1 million of widgets purchased x generic sector-wide spend-factor = X tonnes CO₂”). While easy to implement and audit, the spend-based method is an approximation that is not based on the specific emissions footprint. In other words, it does not identify the actual emissions hotspots, like specific suppliers or factories, and by implication, does not enable targeting transition finance to these hotspots. In Project Symbiosis we investigate whether we can change this through the use of AI for emissions calculation and transition finance matching.

⁸ WBCSD and WRI, Greenhouse Gas Protocol: technical guidance for calculating scope 3 emissions, 2013, ghgprotocol.org/scope-3-calculation-guidance-2.

D. Knock on implications

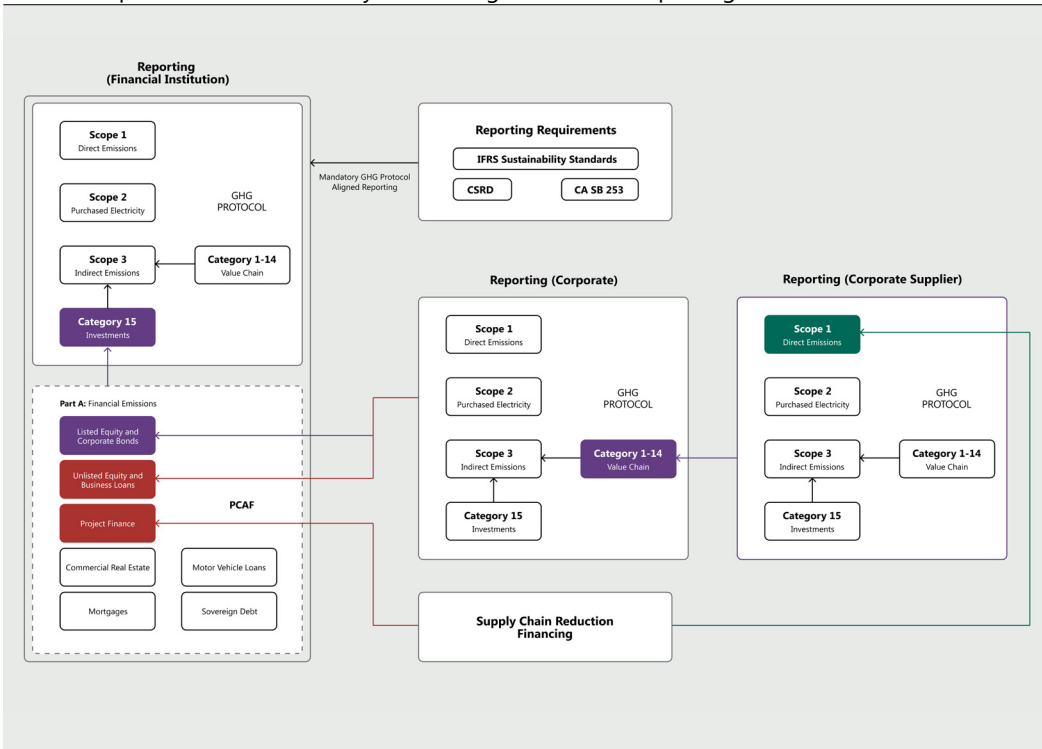
The challenge is even greater within the financial sector, where the vast majority of emissions are concentrated within a single scope 3 category, namely category 15, emissions from investments. To improve methodological consistency, the **Partnership for Carbon Accounting Financials (PCAF)** has adopted a series of standards in recent years against which financial institutions can report more consistently on their investment-related emissions.⁹ Graph 1.2 documents the position of financial institutions within this framework.

However, even in the best-case scenario, such investment emissions are ultimately composed of proportional levels of the reported emissions of their investee companies. Given the difficulty investee companies face in calculating their own scope 3 emissions, financial institutions instead typically rely on economic activity-based calculations for category 15 reporting (ie the spend-based method).

While this approach is very effective for the purposes of consistent accounting, it comes at the expense of visibility into the actual sources of emissions along the supply chain, as such granularity is replaced by top level emissions factors.

Graph 1.2

Relationship between sustainability-related regulations and reporting standards



⁹ Partnership for Carbon Accounting Financials (PCAF), The global GHG accounting and reporting standard for the financial industry: Part A: financed emissions, p 65, (table of calculation approaches for category 15 financed emissions in loans and unlisted equity context), 2023, carbonaccountingfinancials.com/standard#the-global-ghg-accounting-and-reporting-standard-for-the-financial-industry.

E. Impetus to explore the use of AI and related technologies

Breakthrough solutions are needed which utilise a physical activity-based approach that incorporates actual supply chain data to generate emissions calculations (ie which identify actual SME and other upstream sources and the relative magnitude of their emissions impacts).

Such technology solutions, applied at scale, would expose supply chain emissions sources and their comparative intensity, thereby providing a valuable data set for identifying opportunities for emissions reductions. This would enable the application of defined criteria to assess opportunities, eg the alignment of such opportunities with financial profitability metrics, government financing programme requirements or regulatory classification frameworks like the EU taxonomy for sustainable activities.

A scan of AI and related technologies in this area highlights a mixture of common and niche use cases, including many which are within the scope of this project to evaluate. While the real-world effectiveness of these approaches cannot be verified through market scanning alone, it is encouraging to see willingness to adopt and explore how these technologies can address a broad range of problems.

Table 1.1

AI and Related Technologies within the carbon accounting space.

Technology	Use Case	Evaluated in Project Symbiosis
Deep Learning and NLP	Unstructured Data Processing: Extraction and processing of data from unstructured sources such as invoices.	Yes
Deep Learning and NLP	Data Classification: Classification of attributes and characteristics into a common taxonomy.	Yes
Deep Learning and NLP	Climate Risk Analysis: Analysis of policy documents, disclosures and estimated costs to forecast the financial exposure due to climate change.	No ¹⁰
Generative AI, LLM	Decarbonisation Roadmaps: Identification and assessment of possible reduction scenarios.	Yes

¹⁰ Bank for International Settlements, Project Gaia: Exploring tokenised asset markets (2023) <https://www.bis.org/publ/othp84.htm>.

Technology	Use Case	Evaluated in Project Symbiosis
Generative AI, LLM	Methodology Explainer: Explanation of complex or unclear accounting methodologies.	No
Generative AI, LLM	Synthetic Data Creation: Creation of data to support training of climate-related models.	No
Data Science and Analytics	Data Proxy Selection: Selection or inference of missing data points during calculation.	Yes
Data Science and Analytics	Impact Analysis: Analysis and identification of high-impact emissions areas.	Yes
Data Science and Analytics	Emissions Forecasting: Time-series projection of emissions, based on historical data.	No

Glossary

- **Deep learning and natural language processing (NLP)** – deep learning (a subset of machine learning) is a type of AI that enables computers to learn patterns from large amounts of data, similar to how humans learn from experience. NLP is a branch of AI that helps computers understand, interpret and generate human language.
- **Generative AI** – a type of AI that creates new content, such as text, images, music or videos, based on what it has learned from existing data. It powers tools like AI-generated art, chatbots and content writing assistants.
- **Large language models (LLMs)** – a type of AI designed to understand and generate human language. LLMs are trained on vast amounts of text data and can perform a wide range of tasks such as answering questions, summarising content and translating languages. They are a core technology behind many generative AI applications that involve text, such as chatbots and writing assistants. While LLMs are a subset of generative AI, not all generative AI relies on language – some generate images, music or code.

- **Data science and analytics** – the process of collecting, analysing and interpreting data to find patterns, make decisions and solve problems. It helps businesses understand trends, improve operations and predict future outcomes.

III. Project Symbiosis: exploring AI and related technologies for scope 3 accounting and transition finance

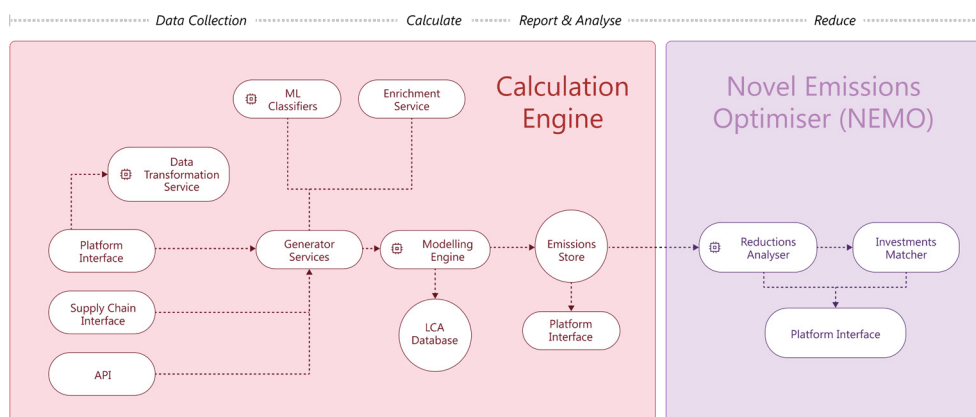
A. Overview

Project Symbiosis explores whether AI offers pathways to replace spend-based with activity-based, supplier-level data with similar ease of implementation.

The first goal of Project Symbiosis is to explore how advanced data techniques and AI can be leveraged to **more accurately collect, interpret and calculate scope 3 emissions** and other impact data in corporate supply chains. For the purpose of this goal, the project performed applied technology research that explores and explains the AI techniques that can be used to achieve this. The second goal of Project Symbiosis is to explore how AI can **identify emissions reduction opportunities** based on such data and thereby lay the groundwork for reducing scope 3 emissions. Finally, Project Symbiosis aims to explore how to leverage emissions and reductions data to **match suppliers with funding sources** to decarbonise the supply chain (referred to as financeable emissions reduction opportunities). For the purpose of these goals, the project developed code referred to as **Novel Emissions Optimiser** or **NEMO**.

Graph 1.3

Architecture of the explored Project Symbiosis platform



Taken together the applied technology research and NEMO provide a blueprint that seeks to address **a number of the challenges** outlined previously, with a specific focus on:

- Reducing the friction and effort involved in **data collection** to calculate a corporate carbon footprint (CCF). High levels of friction consistently result in lower levels of footprint accuracy and a reduced focus on decarbonisation measures.
- Increasing the accuracy of emissions **calculations** by utilising a more granular set of emissions factors and using modelling techniques for targeted proxy selection of missing data points consistent with a range of **reporting** regimes.
- Helping users identify and understand the **reduction potential** of impactful decarbonisation measures, with the ability to forecast and align projected emissions to a target.
- Enabling **users throughout the value chain** to **more accurately** calculate and disclose their emissions to **reduce the compound effect** that can result in poor quality emissions results at later stages of the value chain (ie Scope 3, Category 15).

B. User groups


Project Symbiosis approached these problems by considering the perspective and needs of **multiple user groups within the value chain** (see Graph 1.4):

- **SMEs** must be enabled to calculate the footprint of their own products and business activities in order to disclose the emissions to their own business customers. Although shown as a single user group (Graph 1.4), in practice the supply chain can consist of multiple tiers. As supply chain suppliers are the root of the emitting activity, a focus on trying to maximise the accuracy of the calculations at this stage is vital to minimising the compound effect that layers of data inaccuracies can lead to.
- **Large corporations** must also be enabled to directly calculate the footprint of their products and business activities. As their business models are the drivers of the emissions within their supply chain, and with the regulatory requirements that are coming into force in many developed markets, large corporations are currently a main user group to be served. The majority of the features used by SMEs and large corporations are the same, although large corporations in many markets have stricter requirements around disclosures and reporting.
- **Financial institutions** would not directly calculate and disclose their own footprints through the explored approach, but would instead indirectly benefit if the approaches result in increased accuracy of the disclosures from SMEs and large multinational corporations. With access to such data, financial institutions may be incentivised to take a more active leadership role in driving supply

chain decarbonization through setting lending criteria, prioritising stewardship, or developing new green products.

Graph 1.4

Relationships between user groups across the value chain



C. Explored approaches and outcomes

1. Data collection

Data collection is typically among the most challenging aspects of carbon accounting for a number of reasons. This includes the fragmented nature of supply chains, which may span thousands of supplier entities for large corporate businesses across multiple tiers, where each supplier may store data in different formats and with varying levels of quality and completeness. As these data are the foundation of carbon accounting, approaches were explored to simplify and improve the consistency of data collection.

Methods explored include:

- standardised software-based approaches for supply chain data collection from SMEs that eliminate room for variation in data input;
- application programming interfaces (APIs) to automate the intake of data from corporate systems; and
- automated transformation of the disparate data formats typically provided to corporate entities from diverse suppliers and business units within the corporate entity itself.

The use of AI in connection with these approaches was explored, though performance has been relatively weak to date, as discussed in Part 2. Use of AI to analyse raw input data and then transform it into formats suitable for credible and consistent impact calculations was inconsistent, reflecting the variety and context-specific nature of many data exports.

Misclassifying incoming business data has significant follow-on consequences for calculations and reduction measurement (even minor errors in interpretation can cause cascading ripple effects), meaning high performance is critical for widespread use of AI to improve existing data collection processes. Nonetheless, given the


time-intensive and manual nature of the task, further exploration of AI is warranted as AI technologies continue to rapidly improve.

2. Calculation and reporting

Project Symbiosis next explored approaches for turning real-world data into reliable emissions estimates. We explored a flexible modelling-based approach. This included exploration of a raw data standardisation and enrichment service (generator services) and AI data classifiers to further structure inputs, as well as a modelling engine capable of ingesting such data and automatically computing sufficiently accurate calculations utilising an underlying impact factor database.

Graph 1.5

The explored calculation approach



As discussed in greater detail in Part 2, the modelling engine utilised a Neo4j graph database dual-plane architecture to structure emissions data and harmonise it across activity taxonomies and geographies. It further incorporated advanced statistical modelling techniques, including probability distributions and quantile reconstruction, to infer missing values with high confidence, generating results in alignment with recognised methodologies.

Overall, the explored calculation approach demonstrated potential for the use of AI and related technologies deploying advanced statistical modelling techniques to automate calculations, particularly when there is a requirement for sensitivity to a range of data availability scenarios and calculation at large scale (eg across large corporate data sets).

3. Reduction analysis

Finally, a core focus of Project Symbiosis involved exploring the use of AI for identification of financeable emission reduction opportunities. To accomplish this, we utilised **NEMO**, a **proof of concept software application** developed as part

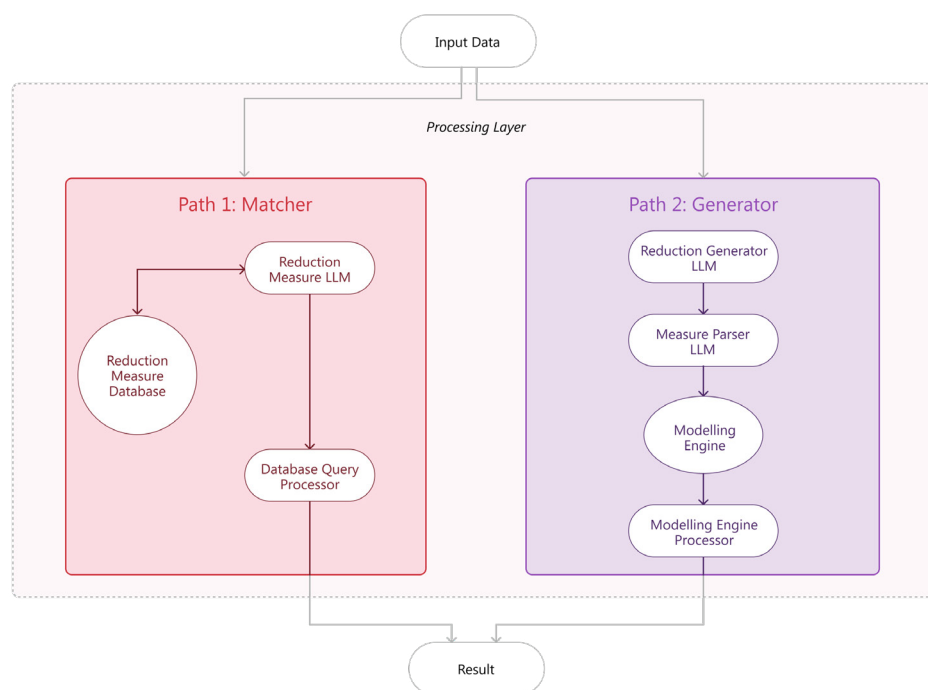
of Project Symbiosis to assess the ability of AI to identify actionable reduction measures based on varying levels of emissions data.

Two alternative AI pathways were explored in NEMO, the **“matcher” path** and the **“generator” path**, each of which deployed LLMs to identify reduction measures through distinct calculation approaches as follow:

- The matcher processing path enabled the selection of reductions from a curated, pre-computed database of reduction measures validated by domain experts. This could be thought of as a top-down approach and was developed to be flexible to a wide range of input data of varying methodological formats.
- The generator processing path, by contrast, sought to dynamically identify and evaluate reduction measures through direct integration with a calculation modelling engine. This could be thought of as a bottoms-up approach and was explored to enable consistent use of the same underlying methodology and emissions factors, enabling a like-for-like comparison of reductions to the baseline.

Graph 1.6

Overview of the processing paths in NEMO



As addressed in depth in Part 2, the processing paths explored through NEMO both showed degrees of promise:

- The matcher path demonstrated high degrees of correctness and minimised hallucinations – an outcome consistent with the curated nature of its underlying reductions data set – yet struggled with target completeness, often returning fewer than half of applicable measures despite prompting and requiring a

high level of dependence on subject matter experts to curate sector-specific measures.

- The generator path demonstrated strong and consistent degrees of accuracy, specificity and applicability, with the development and evaluation process showcasing that targeted prompt engineering – adding prioritisation, real-world constraints and assumption transparency – progressively improved performance. Ultimately, results produced through the generator path approached the quality that might be produced by a human subject matter expert, showcasing the potential of AI to assist in bridging information gaps.

4. Applied example

An example use case for the approaches described in Parts 1 and 2 is provided in Part 3. While the use case is set in the consumer goods context where corporate entities often possess limited, inconsistent primary data on the emissions-generating activities of their supply chain partners, the learnings are intended to be of a generic nature and can be viewed as an illustration of the opportunities and challenges relating to the use of AI to enable calculations and matching. The consumer goods example is also used to illustrate and better explain the technology approaches in Part 2.

Why focus on consumer goods?

AI solutions can be tailored and fine-tuned more optimally if they are specialised for the needs of a specific sector. We chose the consumer goods sector as consumer goods companies (retailers) purchase substantially all of the products they sell to end customers as finished products from third-party manufacturers. As a result of these arm's length relationships, retailers often have very little visibility into their supply chains beyond the contractual relationship with direct suppliers of their finished products. The consumer goods sector's scope 3 emissions have been estimated to generate approximately 20% of all global emissions each year – with the vast majority of that coming from upstream suppliers.¹¹

¹¹ McKinsey & Company, Retailers' climate road map: charting paths to decarbonised value chains, July 2024, www.mckinsey.com/capabilities/sustainability/our-insights/retailers-climate-road-map-charting-paths-to-decarbonized-value-chains#; McKinsey & Company, Tackling Scope 3 emissions through supplier collaboration, January 2024, <https://www.mckinsey.com/capabilities/sustainability/our-insights/sustainability-blog/tackling-scope-3-emissions-through-supplier-collaboration> (together "McKinsey 2024").

IV. Learnings and next steps

Project Symbiosis set out to explore and showcase how novel technologies may offer a viable pathway to addressing significant data gap challenges that hinder improved sustainability and climate performance. It specifically sought to assess novel AI approaches for identifying and calculating emissions throughout a company's extended supply chain, identifying viable reduction opportunities and connecting those opportunities with appropriate financing instruments. This would benefit a wide range of stakeholders across the ecosystem including large multinational corporations, the financial sector, SMEs and ultimately people and the environment.

Significant barriers exist to the robust corporate data collection necessary in order to calculate scope 3 emissions with sufficient detail to generate actionable insights. Further, even as sustainability disclosure grows more standardised and regulated, variations in permissible calculation approaches even within leading standards means that data are rarely comparable across companies. Accordingly, significant expert manual work is required to identify reduction opportunities that are specialised to each company's unique calculation approach. This seriously impedes companies from engaging in such efforts, as well as the ability of financial institutions to fund reduction opportunities at scale.

As demonstrated through Project Symbiosis, AI and related technology approaches have the potential to significantly improve the status quo, whether by improving the speed, breadth and quality of primary data collection, or by generating intelligent environmental impact results that are flexible to a wide range of data availability scenarios. Significant further exploration is warranted about the use of AI to improve and standardise data collection, which is often the most significant barrier to accurate scope 3 calculation.

In terms of generating actionable reduction opportunities, the NEMO application developed for Project Symbiosis demonstrates how AI can streamline the process of assessing reduction opportunities using existing calculated sustainability data through a range of techniques. Some of these techniques approached the expert-level of quality and performance needed to generate an actionable decarbonisation roadmap, although performance varied in unexpected ways across the range of LLM models explored. This demonstrates real potential in the near future for AI to help bridge knowledge gaps that are prevalent even within large companies.

Further exploration of AI may be particularly warranted in connection with the follow-on step of identifying and matching appropriate financing opportunities to those reduction measures, the feasibility of which was explored only at a general level in connection with the project. In particular, the foundational development work underlying NEMO's exploration of AI to identify reduction measures could be utilised and expanded by banks in Hong Kong SAR to encourage SME matched transition finance – a simple and effective way to predict high impact measures likely to require financing.

Importantly, NEMO's performance and development depended heavily on the involvement of subject matter life cycle assessment (LCA) experts – specialising in the economic sector of focus – in the ongoing fine-tuning and refinement process. This is because emissions sources and the viability of reduction opportunities vary significantly across product supply chains. While this suggests an ongoing need for specialist support for tool development and that AI is unlikely to be a one-size-fits-all solution, it does not undermine the overall promise demonstrated by the AI approaches explored and tested in Project Symbiosis.

Additionally, it bears noting that the environmental and climate performance of AI technologies themselves were not explored in connection with Project Symbiosis. These emissions and other environmental impacts such as water consumption may ultimately be substantial, given the significant computing resources required, cooling demand, and the electricity load needed to serve it. While certain LLMs may prove less resource intensive than others, any use of AI is likely to generate significant emissions, even as electricity grids worldwide continue to slowly decarbonise at different paces.

The findings of Project Symbiosis are consistent with and build upon other recent exploratory work on AI and emerging risk undertaken by the BIS Innovation Hub Hong Kong Centre, including Project Gaia's use of AI to extract and assess data from publicly reported corporate climate disclosures. In the future, such data could form the basis for more nuanced, timely and standardised emissions factors for a range of corporate emissions sources that might be utilised in AI-based calculations, which in turn could improve problems with the standardised calculation of scope 3 sources across companies.

As AI developed rapidly over the course of Project Symbiosis and is poised to continue to improve significantly in the near future, further exploration and development of its potential is warranted in connection with the use cases explored. This includes the use of AI to improve supply chain data collection and standardisation, as well as the development of accurate and actionable reduction roadmaps that large multinational corporations, SMEs and financial institutions can rely on to rapidly reduce emissions and other environmental impacts.

Acknowledgements

Project team

*Bank for International Settlements
(BIS) Innovation Hub*

Bénédicte Nolens, Centre Head, Hong Kong Centre

Jean-Marc Champagne, Advisor

Saleh Alqaryan, AI Advisor

Shauna Daswani, Team Coordinator

*Hong Kong Monetary Authority
(HKMA)*

Teresa Lin, Senior Manager (Market Development), External Department

A special thank you to Max Zhou, student of LITE Lab at University of Hong Kong, for his legal background research. We would also like to thank Andrew Chau, Corinne Ho, Henry Holden, Karmela Holtgreve, Francesca Hopwood Road, Kevin Tsui and Eric Wong of the BIS for their review of the report and their feedback. Furthermore, we would like to thank our industry partners for supporting this project and all participants in our engagement and scoping workshop for their participation, including the Bank of East Asia Limited, the Climate Bonds Initiative (CBI), the Carbon Disclosure Project (CDP), the Dubai Financial Services Authority (DFSA), Invesco Plc, The Hong Kong Shanghai Hotels Limited, and Standard Chartered Bank (Hong Kong) Limited.