

**How Can A Computational Model
Guide the Optimisation of
Interdisciplinary Team Design?
- A Pilot Based on Heathrow
Airport's Sustainability Goals**

Executive Summary

Traditional models of knowledge transfer and collaboration, such as social network analysis and bibliometric mapping, focus on relational dynamics and in doing so, overlook the epistemic content of disciplinary knowledge and thus its influence on effective interdisciplinary collaboration. Addressing this gap, this paper presents a novel adaptation of the Knowledge Transfer Graph (KTG), originally developed by Minami et al. (2019), to optimise interdisciplinary team formation for solving complex real-world problems. Using Heathrow Airport's decarbonisation goals as a pilot case, the model assesses the epistemic alignment of distinct disciplinary perspectives, represented by thematic vectors derived from AI generated action places, against a benchmark vector constructed from empirically validated, successful decarbonisation strategies. Through Kullback-Leibler (KL) divergence, minimisation and constrained optimisation, the model quantifies the cognitive value of each discipline as part of an interdisciplinary approach to tackling airport carbon emissions.

Tested, using synthetic benchmarks with known weights, this adaptation of the KTG offers a data-driven and conceptually grounded framework for optimising interdisciplinary team design. It achieves this through evaluating the alignment of an interdisciplinary team's collective knowledge set with the epistemic demands of a complex problem.

The pilot application to Heathrow's sustainability goals demonstrates its practical utility as a guiding framework for interdisciplinary team design in response to complex problems.

Contents

Executive Summary.....	2
Literature Review	4
Bridging Theory and Practice	5
Foundational Knowledge Transfer Graph (KTG) (Minami et al., 2020)	6
Adaption of the KTG	7
Case Study	8
Qualitative Method - Thematic analysis	9
Introduction.....	9
Preparation for analysis.....	9
Thematic Coding Process:.....	10
Limitations of Thematic Analysis	12
Translation to Quantitative Data.....	13
Quantitative Approach – Data-Science, Natural Language Processing (NLP) and Quantitative Modelling.....	13
Introduction to model development	13
Generating Disciplinary Representations.....	14
Kullback-Leibler (KL) Divergence Based Evaluation	14
Individual Alignment – Phase 1	15
Constrained Optimisation of Interdisciplinary Synthesis – Phase 2	15
ASHA-Based Subset Evaluation – Phase 3	16
Limitations.....	16
Methodological Summary.....	16
Results	17
Individual Alignment.....	17
Team Optimisation	18
Team Discovery (ASHA Algorithm)	19
Conclusion	19

Literature Review

Effective knowledge transfer is central to innovation and competitive advantage. It involves the exchange or dissemination of knowledge between individuals, groups or organisations and is closely tied to ‘absorptive capacity’ – the ability to recognise, re-order and apply new information based on prior knowledge (Cohen & Levinthal, 1990). Whilst literature emphasises the importance of acquiring, managing and optimising the economic value of knowledge, enhancing problem-solving capabilities and performance, critics have noted an academic overemphasis on factors that encourage knowledge sharing. The investigation of cooperation, trust and organisational culture/structure dominates consideration of the epistemic output of knowledge sharing and how this can be optimised (Castro & Moreira, 2023; Yeboah, 2023).

Further, due to their narrow focus on the dynamics of knowledge sharing channels, many frameworks favour abstract academic models with limited capacity to inform practice (Ward et al., 2009; Castro & Moreira, 2023; Yeboah, 2023). This focus has ultimately led to little progress being made in suggesting how *strategic* knowledge is created and exchanged (Kowalska-Styczeń et al., 2018). By design, deterministic approaches navigate the challenging complexity of conducting an empirical study in the field (Ward et al., 2009) by simplify complex knowledge systems through assuming uniformity, resulting in less representative models.

To overcome these difficulties, computational models such as agent-based modelling (Appendix A), have been deployed to better map knowledge systems. However, these still largely focus on how structural properties influence knowledge diffusion within organisations (Kowalska-Styczeń et al., 2018), rather than the epistemic value of what is being transferred.

Despite the ‘increasing consensus that real-world policy problems are inherently interdisciplinary and cannot be addressed with knowledge from only a single scientific or academic discipline’ (Newman, 2023), current approaches to modelling such collaboration remain limited in both scope and practical utility (Huutoniemi et al., 2010; Newman, 2023; Hvidtfeldt, 2016; Kiss et al., 2019).

Whilst a wide range of conceptual (e.g., Newman, 2023), computational (e.g., Kowalska-Styczeń et al., 2018), and bibliometric models (e.g., Marchiori & Franco, 2020), have been developed to study collaboration, most of these approaches revolve around the structure of social interaction rather than the structure of knowledge itself (Hvidtfeldt, 2016; Fiore, 2008).

Network analysis techniques, that attempt to map the influence of intellectual structure and relationships (Marchiori & Franco, 2020; Alvarez-Meaza et al., 2020) including: social network analysis (SNA), often used to map collaboration patterns among individuals or institutions by looking at co-authorship, communication or project collaborations (Aboelela et al., 2007; Lungeanu et al., 2014) and bibliometric analysis, deployed to quantify the extent of interdisciplinarity through citation diversity, co-citation patterns and author affiliations (Kiss et al., 2019; Newman, 2023), fail to capture the epistemic value of the knowledge being exchanged (Hvidtfeldt, 2016; Miller et al., 2008). This highlights the need for models that engage directly with the content and compatibility of the knowledge being transferred, rather than the systems

transferring it (Kowalska-Styczeń et al., 2018). These models exist above the limitingly narrow level of abstraction that prevents real-world applicability, whilst residing below the higher-level conceptual frameworks that delve into the epistemology and axiology of multi-, inter- or trans-disciplinarity. These conceptual frameworks, whilst useful for classification lack operational precision and contextual grounding (Newman, 2023) diminishing their ability to inform interdisciplinary team formation and knowledge integration (Morse et al., 2007).

Together, these limitations reflect a broader issue that research surrounding interdisciplinarity and knowledge diffusion frequently fails to consider, the effectiveness and contextual relevance of knowledge components and the subsequent impact on a team or system's collective epistemic capacity (Hvidtfeldt, 2017; Morse et al., 2007). Addressing this issue requires a new approach - one that models the cognitive architecture of disciplinary knowledge sets and their integrational value.

A key motivator for this study, as proposed by Hvidtfeldt (2016), is the notion that effective interdisciplinary collaboration is not simply a function of who interacts, but how well their knowledge components - concepts, methods and assumptions - can be integrated (Hvidtfeldt, 2016; Miller et al., 2008; Huutoniemi et al., 2010). Thus, this paper investigates how disciplinary perspectives can be combined in ways that are epistemically efficient and contextually appropriate, through adaptation of Minami et al.'s Knowledge Transfer Graph (KTG).


The generated model does not aim to replace existing tools like SNA or Bibliometrics but rather, to complement them, by focusing on the integration of diverse disciplinary perspectives. The model offers a conceptually grounded, data-driven method for optimising interdisciplinary team composition to address real-world complex problems. Much of the criticism expressed in this paper, concerns academic frameworks' lack of practical utility. The following section, highlights the potential utility of the model in guiding interdisciplinary team formation to address real-world challenges.

Bridging Theory and Practice

This adaptation of the KTG, goes beyond providing a novel theoretical contribution to academia, delivering practical utility in addressing real-world challenges. As mentioned above, academic methods for mapping interdisciplinary knowledge integration remain limited and further, institutional team formation is largely predicated on job labels or social affiliations. Thus, rarely drawing on the optimal epistemic combination to most effectively address a complex problem.

The model's simplicity, logical structure and clearly derived computational results, aim to encourage an institutional shift in the methodological approach to team design. Rather than inferring contributors' value based on job role, social ties or prior collaboration, it provides a practical tool for gauging the contextual relevance and combinatory value of their disciplinary perspectives, grounding team formation in the cognitive demand of the problem at hand.

Foundational Knowledge Transfer Graph (KTG) (Minami et al., 2019)



Knowledge Transfer Graph
(3 nodes)

Figure 1 – Original KTG (Minami et al., 2019)

Developed in response to the limitations of earlier techniques, such as Knowledge Distillation (Hinton et al., 2015) and Deep Mutual Learning (Zhang et al., 2018), the KTG presents a flexible approach to modelling collaborative learning among deep neural networks. It moves beyond the application of fixed Student-Teacher roles by enabling each model (or “node”) to share knowledge with others, in a way that better reflects the complexity of real-world collaboration where learning often involves multiple sources of influence.

The KTG models collaborative learning as a directed graph (Appendix B) where nodes represent an individual model and edges denote the flow of information between them. What distinguishes the KTG is the ability to regulate the degree and direction of transfer through a system of gating functions (Appendix C), that manipulate the degree of influence of the transfer loss (Appendix D) in the model’s parameter optimisation process. The gradient (Appendix F) of each transfer loss is backpropagated only to the target model, not the source. The total loss observed (Appendix G), represents a model’s own supervised learning and the cumulative influence of its peers, ensuring independent but coordinated learning.

To determine which combinations of nodes, edges and gate types produce the most effective output of a given target node, the KTG framework incorporates a graph optimisation formula - the Asynchronous Successive Halving Algorithm (ASHA) (Appendix H). By exploring a range of graph structures, the ASHA evaluates the ability of each configuration to maximise the accuracy of the target node. It allocates more focus to promising configurations while eliminating weaker ones.

Tests demonstrated that KTGs consistently outperform traditional approaches such as knowledge distillation and deep mutual learning, further noting that increasing the number of

nodes and gate types, improved model performance. This highlights the value of regarding which, and to what extent, nodes contribute to the model's learning. In addition to this, the authors found that KTG configurations optimised on one dataset, generalise well to others, capturing general principles of good collaboration that can be translated to other tasks exhibiting 'Cross Task Robustness' (Minami et al., 2019).

Whilst originally designed for neural networks, the KTG offers a widely applicable conceptual framework. It provides a methodology to explore complex epistemic systems and the integration of sub-systems, such as unique knowledge domains or academic disciplines, within them.

Adaption of the KTG

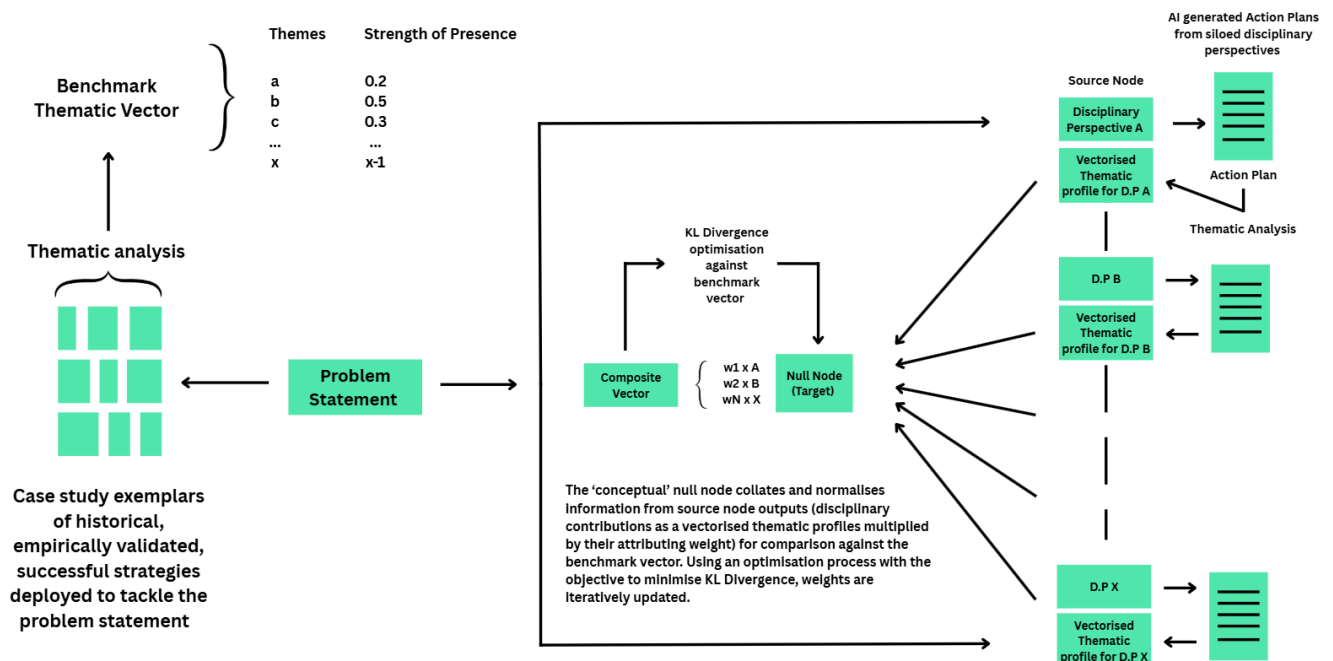


Figure 2 – Reconfiguration of KTG Framework for Interdisciplinary Knowledge Collation

Building on the work of Minami et al., this study recontextualises the core structural logic of the KTG, within an exploratory sequential mixed methods design. The research design aligns with the study's critical realist paradigm, which encourages the use of interpretive qualitative investigation to understand a phenomenon before building a quantitative model to measure it (Hall, 2013). This approach ensures that the adapted model is grounded in real-world data, making it a more robust tool for evaluating how distinct disciplinary perspectives contribute to solving complex, real-world problems.

The adapted architecture illustrated in Figure 2, depicts the relabelling of each node as a disciplinary perspective, connected via a unique edge, to a central null node, that acts as a synthesis point for the collation of disciplinarily siloed insight. The model begins with a problem statement, which guides the collection of a set of contextually relevant benchmark case studies

(peer-reviewed academic papers, official sustainability reports and industry case studies), that detail empirically validated success cases relevant to the problem. Through thematic analysis of these, a codebook (Appendix I) is derived. From which, a unique benchmark vector is formed for each exemplar - a weighted representation of key themes $(a, b, c, \dots, x)$ - aggregated, to act as a proxy for the thematic make-up of an effective, robust and holistic intervention.

The generated codebook, is then applied to AI generated, disciplinarily siloed, action plans (Appendix J), with each plan parsed into a thematic vector that conveys the relative prominence of benchmark-derived themes present in each. These discipline specific vectors, serve as the source node outputs. Unlike the original model, where knowledge is shared between neural agents through backpropagated gradients (Appendix K), this adaptation, introduces a conceptual null node as a consistent target and a theoretical knowledge synthesis point. An optimisation process, using Sequential Least Squares Programming (SLSQP), then calculates the ideal weighting of each disciplinary vector through the minimisation of Kullback-Leibler (KL) divergence (Appendix E) between their combined output and the benchmark. The ASHA (Appendix H), is then deployed to evaluate disciplinary combinations that form the most effective interdisciplinary team of a given size. The model's output, therefore, reveals disciplinary utility (via optimisation) and recommendations for team design (via the ASHA).

Importantly, within this mixed methods approach, generative AI is leveraged as a means to support the model's functionality. However, insight is ultimately drawn from thematic analysis and mathematical computation.

Case Study

Heathrow Airport was selected as the pilot case for this study, due to its scale and organisational complexity. As the UK's largest airport and major global hub, it is the world's second most carbon-intensive airport, reporting over 18 million tonnes of CO₂ emissions in 2023 (Heathrow, 2024; Scott, 2025). Additional emissions include 5,844 tonnes of NO_x and 37 tonnes of PM_{2.5} annually, comparable to the pollution from over 3.2 million cars when combined with other London airports (Gayle, 2024).

As a direct address to its environmental impact, Heathrow's *2.0 Strategy* outlines a sustainable transition toward 'Net Zero Aviation' (Heathrow, 2024) by 2050, through electrified ground operations, the use of sustainable aviation fuel (SAF) and modal shifts in staff/passenger transport, as well as other strategies. As part of their ongoing efforts toward sustainable development, they launched a £30,000 innovation prize through their Sustainability Centre of Excellence, which provides a relevant opportunity to test the model. Rather than proposing direct interventions, the study offers a strategic framework to optimise interdisciplinary team formation, demonstrating the model's value as a consultancy style tool.

Methodology

The study is structured using an exploratory sequential mixed methods design, chosen to ensure that the study produces a robust, qualitatively grounded and empirically validated, tool to address complex, interdisciplinary problems (Newell, 2001). Guided by a critical realist paradigm that accommodates both interpretive and objective forms of inquiry (Hall, 2013), an initial qualitative exploration informs a subsequent quantitative analysis (Tashakkori & Creswell, 2007).

Qualitative Method - Thematic analysis

Introduction

This section outlines the thematic coding process used to support the function of the adapted KTG model. A multi-phased thematic analysis was adopted. An initial, inductive analysis of Heathrow's Sustainability Report 2024 (Heathrow, 2024), formed a contextually grounded guiding codebook (Appendix I) that acted to ensure that the later deductive analysis, of the benchmark papers, remained critically focused and conceptually relevant. Importantly, the thematic coding of these exemplars is not restricted solely to a deductive approach, allowing for novel themes to emerge beyond those that feature in the Heathrow report.

Although the Heathrow case is the focus of this pilot, the methodological structure is designed to be transferable, enabling effective interrogation of other complex problems. The grounding paper is not a prerequisite for an effective thematic approach which could equally deploy a purely inductive methodology to identify themes present in the benchmark corpus. However, it does help to ensure that contextual relevance is maintained and case specific value is derived.

This thematic analysis aims to identify and quantify the presence of key themes within a curated selection of airport decarbonisation strategies. These themes are used to generate degree-scaled thematic vectors, which serve as inputs into a computational model.

Preparation for analysis

Benchmark cases (airport decarbonisation strategies) were gathered using a sampling strategy based on a set of inclusion criteria (Appendix L) - met by seven airports (Appendix M). Source modes included peer-reviewed academic papers, official sustainability reports and industry case studies.

The analysis was conducted through a critical realist lens, acknowledging the realness and quantifiability of carbon emissions and subsequent reduction strategies (ontological realism), but accepting that our understanding of them is contextual and shaped by language, societal influence and interpretation (epistemological relativism). The integrated methodology deployed, is supportive of this approach, as it seeks to interpret patterns in the benchmark sources and then transform them into a quantitative input for a computational model.

As the sole researcher, personal bias and context naturally influence the interpretation of data. In light of this, a multi-phased inductive/deductive approach, alongside the application of Braun and Clarke's (2006) six-phase model for reflexive thematic analysis (Appendix N), helps to

ensure reflexivity and openness to emergent, whilst relevant and focused, insight. The analysis is further guided by Naeem et al.'s (2023) 'A Step-by-Step Process of Thematic Analysis to Develop a Conceptual Model in Qualitative Research', in addition to ThematicAnalysis.net (Braun & Clarke, 2006) which provides a concrete grounding to the methodological approach.

Thematic Coding Process:

Heathrow 2024 Sustainability Report

For the initial deductive component of the thematic analysis, core themes were identified in Heathrow's Sustainability Report 2024 (Heathrow, 2024). The analytical process followed phases 2 to 4 of Braun and Clarke's (2006) thematic analysis framework, systematically coding and grouping data into meaningful patterns before collating them into broader themes. Naeem et al. inspired the quotation driven approach which provided increased methodological traceability in theme derivation.

Phase 1 – Inductive Analysis of the Heathrow Sustainability Report

Keyword Selection

The keywords drawn from the Heathrow report followed Naeem et al.'s *6Rs of keyword selection* (Appendix O). For example, from the statement "We supported the use of 180,000 tonnes of SAF in 2024, with a 92% lifecycle emissions saving and £71 million in incentives" the following keywords were extracted: sustainable aviation fuels (SAF), lifecycle emissions reduction, financial incentives and volume commitment. The selection of key quotations reflects specific actions targeted at addressing carbon emissions.

Coding

Once extracted, these keywords were grouped into codes. For example, terms such as "ultra-rapid EV hubs" and "200kW hubs" were coded under EV charging infrastructure, following Naeem et al.'s *6Rs of coding* (Appendix P). The coding process for the Heathrow Sustainability Report remained entirely inductive in an attempt to identify fundamental codes in aviation sustainability strategies, whilst minimising the influence of pre-existing biases. These codes were then labelled as directly or indirectly impacting carbon emissions.

In total, 14 thematic codes were derived, covering core concepts such as energy systems, transport infrastructure and governance mechanisms - formalised in a thematic codebook.

Theme Development

These codes were then grouped into higher-order themes encapsulating Heathrow's decarbonisation strategy. These themes were designed to be both suitably encompassing, whilst specific enough to have a recognisable set of actions associated with them. For instance, codes related to investment and planning, such as 'Carbon Governance and Planning' and 'ESG Disclosure and Certification', were grouped under the theme Climate-Aligned Investment and Governance, guided by Naeem et al.'s *4Rs of Theming* (Appendix Q). This ensured thematic integrity and relevance to the research question.

The thematic map below (Fig.3) provides an example of how keywords informed code creation, which in turn, informed theme creation. See supporting document 'Thematic Codebook' for the full keyword, code and theme derivation process.




Figure 3 – Thematic Map

Phase 2 – Deductive Analysis of Benchmark Sources

In the second phase of the study, the thematic codebook developed during phase 1 was applied to the set of benchmark documents detailing successful decarbonisation strategies. Initially, keyword identification was based on higher-order themes, but these proved imprecise due to their conceptual breadth. Subsequently, this approach was replaced by the use of the previously established codes, narrowing the conceptual lens. Due to the much larger number of codes than higher-order themes, only those with a direct impact on sustainability were selected for analysis.

To avoid the second phase of thematic analysis being constrained by the codes derived from the Heathrow paper, a brief inductive analysis was conducted to identify additional themes. Keywords and phrases relating to these codes, in addition to those previously derived, established a comprehensive list contained in the resultant codebook. Figure 4 provides a visual representation of this process.




Figure 4 – Deductive Keyword Derivation

Limitations of Thematic Analysis

Whilst this thematic analysis offers a structured method for identifying core themes across empirically validated decarbonisation strategies, several limitations must be acknowledged. First, as the sole researcher, my academic background may have shaped theme identification and interpretation, focusing on technical and institutional narratives rather than behavioural or community-led aspects. Second, the use of the Heathrow report as a foundational document, may have introduced an inherent bias, limiting the considered value of other, nuanced approaches.

Furthermore, whilst the use of Naeem et al.'s pre-defined structure was beneficial in maintaining code consistency, its set parameters may have inhibited thematic scope. Additionally, reduction of the codebook to only include direct interventions, may have inadvertently focused the analysis toward quantifiably supported actions, overlooking those more ethical, systemic or socio-environmental.

Another notable limitation is the inconsistency of source type. Lexical variation between academic publications, industry case studies and official sustainability reports may lead to the underrepresentation of certain themes, not due to absence, but rather terminological difference.

Translation to Quantitative Data

This phase represents a critical point of integration in the mixed methods design, officially known as ‘quantitizing’ (Nzabonimpa, 2018), where qualitative data is converted into a numerical format for computational analysis.

Natural Language Processing techniques (Appendix R) were deployed to measure the relative presence of each theme within the benchmark documents. Frequencies were normalised by document length to produce theme vectors for each source. These were averaged, excluding null values, to produce a benchmark vector (Appendix S) , an empirical proxy for the thematic configuration of a successful airport decarbonisation strategy.

A key limitation to this translation of qualitative to quantitative data, is that themes present in many papers are not weighted more heavily in the final vector. This potentially underrepresents their importance in decarbonisation strategies. However, this approach was maintained to ensure that less present themes in the benchmark papers, were still considered to have an equal potential for impact.

Quantitative Approach – Data-Science, Natural Language Processing (NLP) and Quantitative Modelling

Introduction to model development

This model recontextualises the conceptual application of the Knowledge Transfer Graph (KTG) as proposed by Minami et al., through structural and parameter adaptation, to provide a computational method for evaluating interdisciplinary collaboration. Whereas the original model utilised probabilistic classification outputs from source nodes and minimised Kullback-Leibler (KL) divergence to optimise knowledge transfer, this adaptation replaces classification probabilities with thematic vectors. These vectors are derived from AI generated action plans that represent distinct disciplinary perspectives. KL divergence is maintained as the central optimisation objective (Appendix T), but has been repurposed to quantify the epistemic alignment between each disciplinary vector and the benchmark vector discussed previously.

In place of image classification data, this adaptation relies on thematically coded textual data, used to derive a benchmark vector that mimics the conceptual make-up of a successful decarbonisation strategy.

A prerequisite thematic codebook was created during the qualitative phase of the study, which acts similarly to label taxonomies (Appendix U) in traditional machine learning. Themes were curated based on preexisting accounts of successful carbon reduction strategies and included, for example, SAF ecosystems, electric vehicle transition and building efficiency. Each theme was linked to a corpus of keywords and phrases found in the benchmark texts.

Generating Disciplinary Representations

In the original KTG model, each node operates as an independent learning system, that varies in its architecture and parameters, to produce a unique probabilistic output over a shared input space. In this adaptation, each node is reconceptualised as a distinct academic discipline, instantiated by a GPT-4o-generated "action plan", written in response to a shared prompt. Prompt design was tightly controlled to enforce strict disciplinary boundaries and prevent interdisciplinary outputs, maintaining the modular epistemic architecture of the original KTG design.

Initial implementation involved generating a single plan per discipline. However, early evaluation revealed significant variability in the frequency and weighting of identified themes in each trial due to the stochastic nature of Large Language Models (LLMs). This method was subsequently refined to generate three independent outputs per discipline, mirroring the role of mini-batch sampling in Minami et al.'s neural KTG, where repeated input exposure helps stabilise learning and improve generalisation.

Each plan was then parsed using the thematic codebook to produce its own unique vector (Appendix CC). These discipline-centred vectors replace the softmax-normalised class distributions (Appendix V) used in the original KTG, using distributions across semantically grounded themes instead. Each siloed disciplinary response is represented by a weighted thematic profile, capturing the conceptual make-up of these unique knowledge domains. The adapted method preserves the core principles of independent and diverse node outputs within Minami et al.'s (2019) KTG, whilst adapting its batch-style logic to AI text generation.

Kullback-Leibler (KL) Divergence Based Evaluation

The original KTG deployed Kullback–Leibler (KL) divergence to quantify the transfer loss of knowledge, measuring how effectively a source model's output aligned with, and improved, the target model's predictions. This study adapts that use of KL divergence to evaluate epistemic alignment between disciplinary outputs and a benchmark vector. Rather than comparing probabilistic output vectors, the model compares the normalised frequency vectors of each discipline's thematic make-up ($\hat{V}_i$) against the benchmark vector ($\hat{B}_i$) (Appendix E).

In the adapted model, KL divergence, functions not as a training loss, but as an evaluation metric, serving as a proxy for disciplinary relevance and value. A lower KL score indicates that a given discipline's thematic profile is well-aligned with the benchmark, whilst a higher score suggests thematic misalignment, emulating its role in the original KTG as a measure of deviation from an idealised reference. This metric forms the basis of a three phase evaluation process: phase 1, calculates the KL divergence of each discipline from the benchmark in isolation, measuring the thematic alignment and suggested value of each, in tackling the problem; phase 2, uses constrained optimisation to determine the contributory value of each discipline within a team; phase 3, uses subset evaluation to identify the most effective combination of disciplinary perspectives under team size constraints.

Individual Alignment – Phase 1

Initial analysis evaluated the individual thematic alignment of disciplinary vectors Economics, Environmental Science and Engineering, against the benchmark. Initially, a single action plan instantiated each disciplinary perspective; to enhance robustness, this was later updated to an average of three independently generated plans per discipline. Parsed using the predefined codebook, vectors were derived, normalised and compared to the benchmark. This provided a quantitative reflection of each discipline's value in tackling the problem; the derived thematic profile for each, providing insight into their epistemological offering.

Constrained Optimisation of Interdisciplinary Synthesis – Phase 2

In the original KTG, the target node aggregates backpropagated gradients (Appendix W) from source nodes, with diffusion controlled by gate functions and learned weights. These gradients, reflecting divergence between source and target distributions, are used to iteratively adjust model parameters, enabling the system to learn the optimal configuration and extent of source input.

In this adaptation, the goal is not to train a predictive model, but rather to synthesise disciplinary knowledge in a way that most closely approximates a validated benchmark. Each discipline's thematic vector ($\hat{V}_i$), is treated as a potential contributor to a composite interdisciplinary vector ($\hat{V}_w$), formed as a weighted sum of disciplinary inputs (Appendix X). Optimisation then works to minimise KL divergence between the weighted composite vector ($\hat{V}_w$) and the benchmark vector ($\hat{B}$), by manipulating the weights assigned to each disciplinary vector ($\hat{V}_i$) (Appendix Y). This optimisation process is more closely aligned with the optimisation of gating functions in the original paper, where the influence of each node is modulated to produce a more effective result.

Sequential Least Squares Programming (SLSQP) (Appendix Z) is used to solve the constrained minimisation problem. This shift is a key methodological adaptation from the original KTG, moving from gradient-based learning to direct optimisation. Whereas the original KTG optimised internal model parameters to align predictions with guidance, this adaptation optimises the disciplinary make-up of a composite interdisciplinary output ($\hat{V}_w$), such that it best approximates a benchmark.

The resulting weights (w^*) illustrate the optimal blend of disciplinary contributions required for the composite, interdisciplinary output, to most closely reflect the benchmark. A higher weight (w_i), implies that the epistemological offering of discipline(i) plays a larger role in aligning the combined output of an interdisciplinary proposal with an effective, robust and holistic approach to tackle the given problem. Conversely, low or zero weights, suggest a lesser role or redundancy of a disciplinary perspective's contribution to the collective output.

ASHA-Based Subset Evaluation – Phase 3

Whilst the KL divergence-based optimisation determines the ideal weighting of each discipline's epistemological offering to an effective response to the problem, it operates under the assumption that all disciplinary perspectives are simultaneously available and contributing. However, in real-world scenarios, constraints such as limited team size or financial cost make full participation infeasible.

To address this, rather than optimising a single, pre-defined set of disciplines, the Asynchronous Successive Halving Algorithm (ASHA) is used to identify the most effective interdisciplinary team from a larger pool of candidates. This approach evaluates many fixed-size subsets, maintaining the use of KL divergence to evaluate the output of the collective output of a given team. This methodological progression elevates the analysis beyond merely suggesting **how** to structure a team based on **what** each discipline offers, to suggesting **who** should be chosen to form the most effective and epistemically complementary team.

Limitations

Although this model retains many of the core components of Minami et al.'s original KTG, through adaptation, several limitations arise:

- **Quantification of knowledge:** The thematic vectors designed to mimic softmax distributions are derived from keyword counts which overlook nuances in the theoretical and conceptual nature of each disciplines' approach.
- **Subjective Coding:** Unlike traditional models that use fixed categorical labels, themes are defined by researcher-selected keywords, introducing an element of interpretive bias into the quantitative metrics.
- **Restrictive Source Material:** Disciplinary plans are generated by ChatGPT, they are thus unlikely to reflect the entire contributory ability of each discipline.
- **Benchmark Anchoring:** Optimising for similarity to historical success may underrepresent novel disciplinary contributions.

See accompanying notebook for code specific limitations and future directions.

Methodological Summary

This methodology recontextualises the KTG framework within a novel setting. Qualitative disciplinary outputs are translated into normalised thematic distributions, with KL divergence retained as the central measure of epistemic alignment. Through constrained optimisation and ASHA-based subset evaluation, the model identifies the value of disciplines and optimal

configurations of interdisciplinary teams to tackle a complex problem. Through direct adaptation of Minami et al.’s design, this model looks at the output of interdisciplinary collaboration, framing epistemic value as a process of synthesising conceptual, relevant insight, from diverse disciplinary perspectives. This computationally grounded methodology offers a multi-layered blueprint for interdisciplinary team design, addressing **what** each disciplinary perspective offers, **how** they should be combined, and **who** should ultimately constitute the team in order to optimise epistemological output.

Results

The following results are taken from a single run of the model (see supporting document ‘Results of Trial’). Based on the heterogeneity of Large Language Model (LLM) outputs, multiple runs will naturally exhibit variation in results. Measures have been put in place to minimise the degree of variation, with future development focused on enhancing the model’s fidelity and trial consistency.

Individual Alignment

Table 1:

Individual KL Divergence:

Discipline	KL Divergence vs Benchmark*
Economics	0.7044
Environmental Science	0.3819
Engineering	0.6943

*initial analysis measured the thematic alignment of each discipline, in isolation, with the benchmark vector.

With the lowest KL score, **Environmental Science** exhibits a broader perspective that mirrors the thematic composition of a successful approach. This suggests that an individual with a knowledge set rooted in this discipline would be best placed to provide the most effective approach to Heathrow’s sustainability goals when operating in isolation.

Conversely, **Economics** and **Engineering** show high divergence scores. This indicates an unbalanced thematic profile with disproportionate focus on a few critical themes. Whilst the balance of their thematic composition fails to reflect that of the benchmark, their value lies in

their specialised contribution. These distinct profiles reveal the role each disciplinary perspective plays in the development of an effective response to a complex problem.

Team Optimisation

The optimisation process shifts the focus away from individual disciplines, to their combined value within an interdisciplinary team. It identifies how varied disciplinary profiles are best leveraged in a collective response.

Table 2:

Optimised Weights:

Discipline	Weight
Economics	0.6009
Environmental Science	0.1734
Engineering	0.2257

These results provide a clear and efficient team-building strategy that, in this case, prioritises depth of insight from certain disciplines, drawing on the more holistic approach of others, to enhance the solution.

The model attributes the largest weight to **Economics (60.1%)**, utilising its in-depth coverage of the benchmark's most dominant themes such as 'Emission Reduction', to deliver a high-impact contribution. The optimisation result suggests the incorporation of **Engineering (22.6%)** as a complementary discipline to cover essential technical and infrastructural needs.

While the low weight of **Environmental Science (17.3%)** appears paradoxical given its close alignment with the benchmark in isolation, its role is strategic. Its broad thematic profile fills the remaining thematic gaps, creating a final collective profile that is more closely aligned to the benchmark.

To test the validity of the model’s computational approach, a multi-dimensional test was deployed, using a synthetic benchmark with a known ground truth (Appendix AA). The test confirmed that the optimisation process could accurately recover the predefined 'true weights' used to construct the benchmark, in addition to validating the model’s ability to significantly outperform other approaches, that use uniform and random weightings.

Team Discovery (ASHA Algorithm)

To extend the optimisation process, the Asynchronous Successive Halving Algorithm (ASHA) was implemented to evaluate the performance of teams, restricted by size.

A crucial insight from this analysis is that a greater number of disciplines does not guarantee a more optimal team. Overlooking latent value, the model treats strategic complementarity as more critical than the number of contributors, focussing on the additive value of each perspective. Thus, the model minimises the weight assigned to a discipline if it fails to improve the team's initial interdisciplinary, epistemic composition.

Table 3:

Top-performing combinations consistently featured **Engineering** and **Political Science**, with variations in the third discipline:

Top Combination	KL Divergence	Weight Distribution
(Engineering, Political Science, Anthropology)	0.2251	[0.643, 0.344, 0.013]
(Economics, Engineering, Architecture):	0.2392	[0.487, 0.514, 0.000]
(Economics, Environmental Science, Anthropology)	0.2975	[0.682, 0.287, 0.031]

Note: teams were derived from eight disciplines (Appendix BB) and scored based on the KL divergence between their thematic composition and the benchmark.

The model frequently pairs, high-impact, quantitative-leaning disciplines (Engineering and Economics) with those offering a social or humanistic perspective. This encourages a strategy of assembling a predominantly technical core, guided by broader contextual insight when under an $N = 3$ size constraint and limited disciplinary availability. The results ultimately highlight the value of using the ASHA to reflect realistic constraints of team size and personnel availability.

Conclusion

By shifting the academic focus away from the dynamics of collaboration and toward the epistemic make-up of its output, this model provides an innovative, data-driven tool for interdisciplinary team design, reframing it as a strategic synthesis disciplinary perspectives.

The exploratory sequential mixed method design, drawing on qualitative thematic analysis to enable quantitative formulation, conduces a grounded, systematic comparison and optimisation process.

Findings from the initial SLSQP optimisation indicated that disciplinary value doesn't lie solely in individual thematic alignment, but in the complementarity of disciplines within a combined

team. For example, in a team comprising Economics, Environmental Science and Engineering, Economics emerged as the most valuable contributory discipline despite Environmental Science achieving the lowest standalone KL divergence. This highlights the model's capacity to identify strategic value beyond assumed value. The use of the ASHA, to explore a wider pool of disciplines for team formation restricted by size, concluded that high-performing configurations often included a focus on technical domains.

Importantly, this model does not rely on the unrealistic notion that experts are domain specific, but rather, that individuals typically bring with them a dominant disciplinary lens, shaped by their education and experience. The model therefore serves as a framework for identifying the optimal combination of disciplinary knowledge sets within a team, regardless of whether individuals formally identify as disciplinary experts.

The adapted KTG is a generic framework designed to be widely applicable to contexts beyond airport decarbonisation. Using of alternative benchmark sources and source node outputs, its architecture can be applied to any complex problem that demands the integration of disciplinary perspectives, such as healthcare, urban resilience, politics and governance.

Future developments:

- **Benchmark and Batch Expansion:** Incorporating successful decarbonisation reduction strategies of other airports or similar organisation, will help to improve the validity of the benchmark vector. Increasing the size of source node batches will similarly improve their ability to accurately reflect the thematic make-up of a disciplinarily rooted approach.
- **Automation:** Continued development of the model looks to further leverage AI to gather relevant data sources and adapt source node output based on the given problem. This, alongside the implementation of computational-led thematic analysis using NLP, will enable the model to provide guidance on interdisciplinary team formation without human input.
- **Disciplinary Dialogue Simulation:** Future development, drawing on more of the original KTG methodology, could simulate interaction between disciplines via iterative prompting to incorporate intra-node learning – mimicking team collaboration. The model could subsequently be used in place of a human team, effectively collating the most relevant information from the best suited disciplines to provide a solution, validated by previous success cases.
- **Knowledge diffusion consideration:** Introducing factors that influence the flow of information between parties or nodes will help to improve the model's real-world applicability and provide a more holistic appraisal of interdisciplinary collaboration.

Ultimately, this paper delivers more than a theoretical model; it provides a practical, data-driven tool to guide interdisciplinary team formation based on the epistemological approach of disparate disciplinary perspectives. The adapted KTG model, grounded in a mixed methods approach, offers an adaptable framework for strategic interdisciplinary team design that aims to encourage an institutional shift in the current approach that relies on tradition, intuition or job label, to one that is strategically optimised. The model adds value by identifying **what** each discipline has to offer, **how** this insight is best leveraged and **who** ultimately constitutes an optimal interdisciplinary team.

References

- Aboelela, S., Merrill, J., Carley, K., & Larson, E. (2007). Social Network Analysis to Evaluate an Interdisciplinary Research Center The Challenge of Integration. *The Journal of Research Administration*, XXXVIII(1), 61.
- Alvarez-Meaza, I., Pikatza-Gorrotxategi, N., & Rio-Belver, R. M. (2020). Knowledge Sharing and Transfer in an Open Innovation Context: Mapping Scientific Evolution. *Journal of Open Innovation: Technology, Market, and Complexity*, 6(4), 186. <https://doi.org/10.3390/joitmc6040186>
- Braun, V., & Clarke, V. (2006). *Thematic Analysis | Just Another University of Auckland Blogs Sites Site*. Wwww.thematicanalysis.net. <https://www.thematicanalysis.net/>
- Castro, R., & Moreira, A. C. (2023). Mapping Internal Knowledge Transfers in Multinational Corporations. *Administrative Sciences*, 13(1), 16. <https://doi.org/10.3390/admsci13010016>
- Cohen, W., & Levinthal, D. (1990). Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, 35(1), 128–152.
- Fiore, S. M. (2008). Interdisciplinarity as Teamwork. *Small Group Research*, 39(3), 251–277. <https://doi.org/10.1177/1046496408317797>
- Gayle, D. (2024, February 27). London is most exposed city in world to air pollution from aviation, study finds. *The Guardian*. <https://www.theguardian.com/environment/2024/feb/27/london-is-city-most-exposed-to-air-pollution-from-aviation-global-study-finds>
- Hall, R. F. (2013). Mixed Methods: In search of a paradigm. In *Conducting Research in a Changing and Challenging World* (pp. 71–78). Nova Science Publishers Inc.
- Heathrow. (2024, March). *HEATHROW'S SUSTAINABILITY REPORT 2024* . Heathrow Airport. <https://www.heathrow.com/company/about-heathrow/heathrow-2-0-sustainability-strategy>
- Hinton, G., Vinyals, O., & Dean, J. (2015). *Distilling the Knowledge in a Neural Network*.
- Huutoniemi, K., Klein, J. T., Bruun, H., & Hukkinen, J. (2010). Analyzing interdisciplinarity: Typology and indicators. *Research Policy*, 39(1), 79–88. <https://doi.org/10.1016/j.respol.2009.09.011>
- Hvidtfeldt, R. (2016). Interdisciplinarity as Hybrid Modeling. *Journal for General Philosophy of Science*, 48(1), 35–57. <https://doi.org/10.1007/s10838-016-9344-x>
- Kiss, A., Fritz, P., Lakner, Z., & Soós, S. (2019). Linking the dimensions of policy-related research on obesity: a hybrid mapping with multicluster topics and interdisciplinarity maps. *Scientometrics*, 122(1), 159–213. <https://doi.org/10.1007/s11192-019-03293-8>
- Kowalska-Styczeń, A., Malarz, K., & Paradowski, K. (2018). Model of Knowledge Transfer Within an Organisation. *Journal of Artificial Societies and Social Simulation*, 21(2). <https://doi.org/10.18564/jasss.3659>
- Lungeanu, A., Huang, Y., & Contractor, N. S. (2014). Understanding the assembly of interdisciplinary teams and its impact on performance. *Journal of Informetrics*, 8(1), 59–70. <https://doi.org/10.1016/j.joi.2013.10.006>

- Marchiori, D., & Franco, M. (2020). Knowledge transfer in the context of inter-organizational networks: Foundations and intellectual structures. *Journal of Innovation & Knowledge*, 5(2), 130–139. <https://doi.org/10.1016/j.jik.2019.02.001>
- Miller, T., Baird, T., Littlefield, C., Kofinas, G., Chapin, F., & Redman, C. (2008). *Epistemological Pluralism: Reorganizing Interdisciplinary Epistemological Pluralism: Reorganizing Interdisciplinary Research Research*.
- Minami, S., Hirakawa, T., Yamashita, T., & Fujiyoshi, H. (2019). *Knowledge Transfer Graph for Deep Collaborative Learning*.
- Morse, W. C., Nielsen-Pincus, M., Force, J. E., & Wulforth, J. D. (2007). Bridges and Barriers to Developing and Conducting Interdisciplinary Graduate-Student Team Research. *Ecology and Society*, 12(2). <https://doi.org/10.5751/es-02082-120208>
- Naeem, M., Ozuem, W., Howell, K., & Ranfagni, S. (2023). A Step-by-Step Process of Thematic Analysis to Develop a Conceptual Model in Qualitative Research. *International Journal of Qualitative Methods*, 22. <https://doi.org/10.1177/16094069231205789>
- Newman, J. (2023). Promoting Interdisciplinary Research Collaboration: A Systematic Review, a Critical Literature Review, and a Pathway Forward. *Social Epistemology*, 38(2), 135–151. <https://doi.org/10.1080/02691728.2023.2172694>
- Newell, W. H. (2001). A Theory of Interdisciplinary Studies. *Issues in Integrative Studies*, 19, 1–25
- Nzabonimpa, J. P. (2018). Quantitizing and qualitizing (im-)possibilities in mixed methods research. *Methodological Innovations*, pages 1–16. <https://doi.org/10.1177/2059799118789021>
- Scott, A. (2025, February 20). *Heathrow Airport has a plan to reduce its nature impacts. It also plans to add a new runway*. Trellis. <https://trellis.net/article/heathrow-airport-plan-reduce-nature-impacts-new-runway/>
- Tashakkori, A., & Creswell, J. W. (2007). The new era of mixed methods. *Journal of Mixed Methods Research*, 1(1), 3–7. <https://doi.org/10.1177/1558689806293042>
- Ward, V., House, A., & Hamer, S. (2009). Developing a Framework for Transferring Knowledge Into Action: A Thematic Analysis of the Literature. *Journal of Health Services Research & Policy*, 14(3), 156–164. <https://doi.org/10.1258/jhsrp.2009.008120>
- Yeboah, A. (2023). Knowledge sharing in organization: A systematic review. *Cogent Business & Management*, 10(1). <https://doi.org/10.1080/23311975.2023.2195027>
- Zhang, H., Cisse, M., Dauphin, Y. N., & Lopez-Paz, D. (2018). *mixup: Beyond empirical risk minimization*. In International Conference on Learning Representations (ICLR)

This work contains [6000] words