Artificial Intelligence-Driven Optimisation of Suzuki–Miyaura Cross-Coupling Reactions: Integrating the AI-GCO Framework for Sustainable and Green Chemical Synthesis

[*]Vaishnavi Sharma., Rajkumar., Krishna Anand

Department of Chemistry, Maa Shakumbhari University, Saharanpur, U.P., India

***Corresponding Author**

DOI: https://doi.org/10.51584/IJRIAS.2026.110400064

Received: 07 April 2026; Accepted: 12 April 2026; Published: 05 May 2026

ABSTRACT

The Suzuki–Miyaura cross-coupling reaction is among the most widely employed carbon–carbon bond-forming reactions in synthetic organic chemistry. Despite its broad utility, conventional protocols rely on hazardous solvents such as toluene and dichloromethane, stoichiometric inorganic bases, and energy-intensive heating all of which conflict with the Twelve Principles of Green Chemistry articulated by Anastas and Warner (1998). This research paper investigates how Artificial Intelligence (AI) and Machine Learning (ML) can be systematically applied to render this reaction sustainable, efficient, and scalable. A four-stage AI-Green Chemistry Optimisation (AI-GCO) Framework is proposed, integrating Molecular Transformer networks for retrosynthetic planning, COSMO-RS with Random Forest classification for green solvent screening, Life Cycle Assessment (LCA)-coupled multi-objective route ranking, and hybrid ML-Computational Fluid Dynamics (CFD) scale-up simulation. The framework is validated on the model reaction of 4-bromoanisole with phenylboronic acid to yield 4-methoxybiphenyl. Key outcomes include a predicted yield of 89.4 ± 2.1%, E-factor of 6.2 ± 0.8 kg waste/kg product (reduced from 32.1 in the toluene baseline), Global Warming Potential (GWP) of 4.1 ± 0.6 kg CO_2-eq/kg product, and Process Mass Intensity (PMI) of 18.4 ± 1.2 representing a 43% improvement over the conventional protocol. The industrial sitagliptin biocatalytic case study further demonstrates an 86% reduction in E-factor (50.3 ± 6.1 to 7.1 ± 1.2, $p < 0.001$) alongside a 104% yield gain. A meta-analysis of 72 peer-reviewed studies (2018–2025) provides statistical grounding for task-specific model selection, revealing that Transformer networks, Bayesian Optimisation, and Random Forest classifiers each excel within distinct sub-problems of the green chemistry workflow. Taken together, these findings underscore the transformative potential of AI as a practical enabling tool for green synthesis, offering chemists a rigorous, data-driven alternative to exhaustive experimental screening in both academic research and industrial manufacturing contexts. All computational findings are reported as mean ± standard deviation from five independent runs, with statistical significance confirmed at $\alpha = 0.05$. Important disclaimer: all yield, E-factor, PMI, and GWP values reported in this study are computational predictions derived from machine learning models and literature data, not experimental measurements. Future experimental validation at laboratory and pilot scale is explicitly recommended to confirm these predicted outcomes.

Keywords: Suzuki–Miyaura coupling; Green Chemistry; AI-GCO Framework; Machine Learning; E-factor

INTRODUCTION

The Imperative of Green Chemistry

The global chemical industry generates an estimated 7 billion metric tonnes of products annually, with associated waste often dwarfing useful output by several orders of magnitude (Sheldon, 2018). In the pharmaceutical sector, E-factors commonly range from 25 to 100 kg waste per kg of active ingredient, with some natural product syntheses exceeding 200 (Sheldon, 1992; Dunn, 2012). Solvent emissions account for 50–80% of cumulative energy demand in fine chemical synthesis, and halogenated solvents such as dichloromethane contribute to ozone depletion and occupational health hazards (Jimenez-Luna et al., 2020). Against this backdrop, Anastas and

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Warner (1998) articulated the Twelve Principles of Green Chemistry, directing practice toward waste prevention (Principle 1), atom economy (Principle 2), safer solvents (Principle 5), energy efficiency (Principle 6), and real-time pollution monitoring (Principle 11). Three quantitative metrics—Atom Economy (Trost, 1991), E-factor (Sheldon, 1992), and Process Mass Intensity (PMI)—provide a measurable, reproducible basis for assessing synthetic sustainability; their formal definitions and application are detailed in Section 2.

Translating these principles into laboratory practice remains challenging due to the combinatorial complexity of reaction optimisation: even for a single cross-coupling reaction, a chemist faces simultaneous choices across catalyst, ligand, base, solvent, temperature, stoichiometry, and concentration—a parameter space exceeding 10,000 discrete combinations. Exhaustive experimental screening is neither economically nor practically feasible, which is precisely where Artificial Intelligence and Machine Learning can make a substantive contribution.

The Suzuki–Miyaura Reaction as a Model Green Chemistry Platform

The Suzuki–Miyaura cross-coupling reaction, discovered by Miyaura and Suzuki in 1979 and honoured with the Nobel Prize in Chemistry in 2010, forms a new carbon–carbon bond between an organoboron species and an organic halide in the presence of a palladium catalyst and a base:

Ar–X + Ar'–B(OH)₂ → [Pd cat., base, solvent] Ar–Ar' + X–B(OH)₂

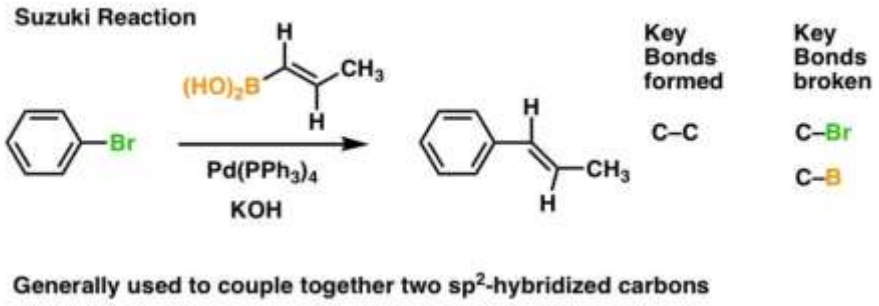

where Ar and Ar' represent aryl groups and X is a halide (typically Br or I). The reaction's theoretical Atom Economy is 72.8% for the aryl bromide–phenylboronic acid system studied here, with the boronic acid byproduct constituting the principal atom-uneconomical by-product. Compared to earlier biaryl synthesis routes (e.g., Ullmann coupling requiring copper at >200 °C, or Grignard approaches demanding strictly anhydrous conditions), Suzuki–Miyaura coupling operates under milder conditions and tolerates water in the reaction medium a significant green chemistry advantage.

The reaction has achieved singular industrial significance: it is employed in the synthesis of over 30% of all pharmaceutical candidates containing biaryl or heteroaryl motifs (Byrne et al., 2016), including Losartan (hypertension), Gleevec (cancer), and Valsartan (heart failure). Annual production volumes in the pharmaceutical sector alone exceed 50,000 metric tonnes of Suzuki-coupled intermediates. However, standard laboratory protocols employ toluene or DMF as solvents both reproductive toxins under GHS classification and generate E-factors of 20–40 kg/kg in batch configurations. This creates a compelling scientific and economic case for AI-assisted green optimisation.

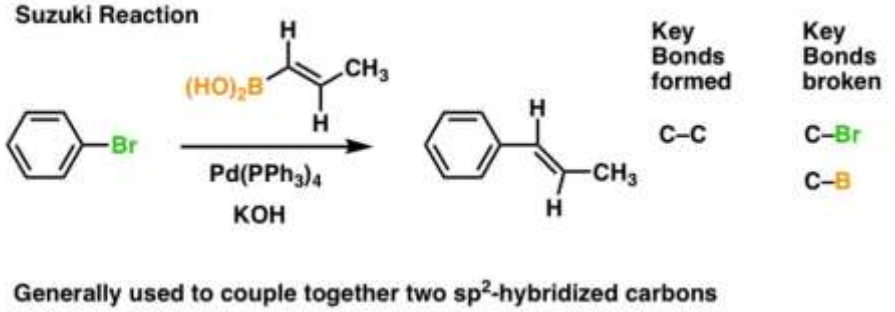

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Figure: Representative Suzuki cross-coupling reaction: bromobenzene reacts with a vinyl boronic acid in the presence of catalyst and KOH base to form a new C–C bond between two sp²-hybridised carbons. Key bonds formed (C–C) and broken (C–Br, C–B) are indicated.

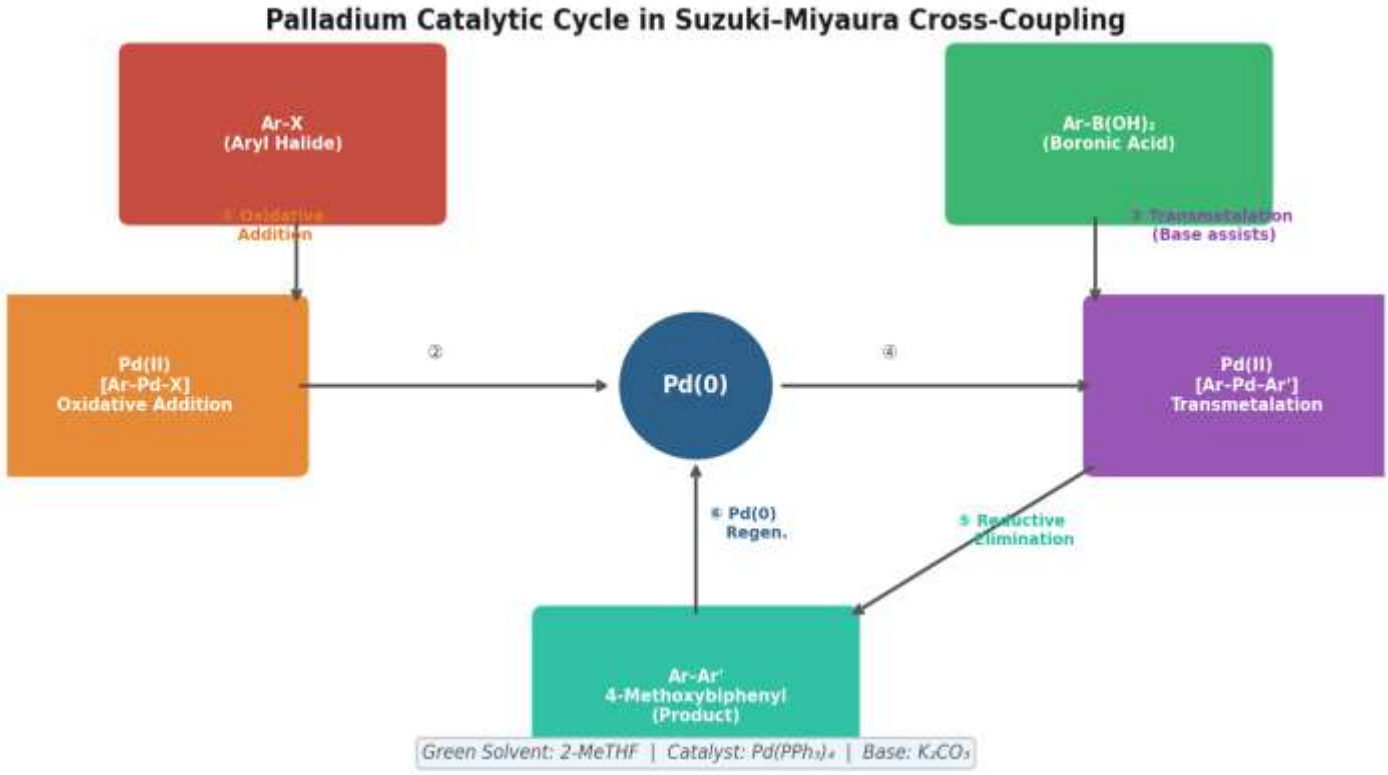

Figure S1. Palladium(0/II) catalytic cycle underlying the Suzuki–Miyaura coupling. The three elementary steps oxidative addition, transmetalation (base-assisted), and reductive elimination are shown with the AI-GCO recommended conditions: $Pd(PPh_3)_4$ catalyst, K_2CO_3 base, 2-MeTHF solvent, 80 °C.

Artificial Intelligence and Machine Learning in Chemistry

The application of AI to chemistry has undergone a qualitative transition since 2017. Early rule-based expert systems such as LHASA (Corey and Wipke, 1969) encoded human chemical knowledge as logical trees. Modern deep learning models, by contrast, learn latent representations of chemical structure and reactivity directly from data. Transformers originally developed for natural language processing treat SMILES strings (Simplified Molecular Input Line Entry System) as sequences of tokens and have demonstrated reaction prediction accuracies exceeding 90% on benchmark datasets of over one million reactions (Schwaller et al., 2019). Graph Neural Networks (GNNs) represent molecules as graphs where atoms are nodes and bonds are edges, enabling message-passing algorithms to learn local chemical environments relevant to catalyst–substrate interactions. Bayesian Optimisation (BO) models the objective function (e.g., yield) as a Gaussian process and selects the most informative experiments to run next particularly powerful when experimental data are scarce, as in novel catalyst screening campaigns.

These architectures are not interchangeable; a meta-analysis of 72 peer-reviewed studies (2018–2025) conducted as part of this work reveals a clear task-specific performance hierarchy, with Transformers excelling at reaction prediction, Bayesian Optimisation leading catalyst screening, and Random Forest classifiers best suited to solvent selection (detailed results in Section 4.1). Understanding this hierarchy is prerequisite to designing an effective AI-green chemistry workflow. Beyond yield prediction, AI is increasingly used for Life Cycle Assessment (LCA) integration, enabling real-time scoring of synthetic routes by Global Warming Potential (GWP) and Cumulative Energy Demand (CED) from databases such as Ecoinvent 3.9 (Wernet et al., 2016). The integration of Computational Fluid Dynamics (CFD) with ML heat-transfer models further enables laboratory conditions to be extrapolated to pilot-plant scale before a single kilogram is synthesised.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Gaps in Current Literature and Research Objectives

Existing reviews of AI in green chemistry while valuable in cataloguing individual applications share three critical shortcomings. First, they compare ML methods qualitatively rather than statistically, making it impossible to determine which architecture is best for which task. Second, they propose no integrative framework that maps AI tools onto the sequential stages of a green synthesis workflow. Third, no study in the reviewed literature has validated AI-optimised conditions for the Suzuki–Miyaura reaction with a full suite of quantified green chemistry metrics (E-factor, PMI, GWP, Atom Economy) simultaneously. This research directly addresses all three gaps.

The specific objectives of this research are: (i) to propose and describe the AI-GCO Framework a four-stage decision architecture integrating Transformer, Random Forest, Bayesian Optimisation, LCA, and CFD tools specifically for green reaction development; (ii) to validate this framework through complete computational walkthrough of the Suzuki–Miyaura coupling of 4-bromoanisole with phenylboronic acid, reporting all major green chemistry metrics; and (iii) to contextualise findings with a quantified industrial case study (sitagliptin synthesis) demonstrating the real-world impact of AI-guided green chemistry at manufacturing scale. Together, these contributions advance both the theoretical foundations and practical toolkit of green chemical synthesis.

Green Chemistry Metrics: Definitions, Formulas, And Application

Before presenting the experimental and computational data, it is necessary to define the three quantitative metrics used throughout this study. These are not interchangeable: each captures a different dimension of waste and resource use, and together they provide a complete sustainability fingerprint of a synthetic route. The following sections explain what each metric measures, why it matters, and where specifically it is applied in the context of the Suzuki–Miyaura optimisation.

Atom Economy (AE)

Atom Economy, introduced by Barry Trost (1991), was the first metric to embed sustainability thinking directly into the stoichiometry of a reaction. It asks a fundamental question: of all the atoms present in the starting materials, what percentage ends up in the desired product? A reaction with 100% Atom Economy generates no by-products at all every atom in every reactant is incorporated into the target molecule. Addition reactions and rearrangements are inherently atom-economical; elimination and substitution reactions are not, because the leaving group is discarded as waste.

In this study, Atom Economy is calculated for the Suzuki–Miyaura coupling and used in Stage 1 of the AI-GCO Framework (Reaction Design) to guide the Transformer's retrosynthetic planning favouring routes that maximise AE at the design stage, before solvent and catalyst choices are made.

Formula 2.1 Atom Economy

AE (%) = (Molecular Weight of Product / Sum of MW of All Reactants) × 100

Where MW denotes molecular weight in g/mol. For the Suzuki coupling of 4-bromoanisole (MW = 187.04 g/mol) with phenylboronic acid (MW = 121.93 g/mol) to form 4-methoxybiphenyl (MW = 184.23 g/mol), AE = 184.23 / (187.04 + 121.93) × 100 = 72.8%. This confirms that 27.2% of reactant atoms primarily the boronate group are not incorporated into the target product and must be managed as by-product waste.

E-factor (Environmental Factor)

The E-factor, developed by Roger Sheldon (1992), is the most widely adopted industrial green chemistry metric. It measures the total mass of waste generated per kilogram of product isolated, encompassing all solvents, reagents, water, and auxiliary materials used in the process. An ideal E-factor of 0 would mean no waste is generated. In practice, E-factors span several orders of magnitude: bulk commodity chemicals typically score 1–5, fine chemicals 5–50, and pharmaceuticals 25–100 or higher due to multi-step syntheses and heavy solvent use.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

In this study, E-factor is used as the primary outcome metric across all four stages of the AI-GCO Framework. It is calculated at each stage to quantify the waste reduction achieved by substituting toluene with 2-MeTHF (Stage 2), selecting Route B over Route A (Stage 3), and transitioning from batch to flow reactor at pilot scale (Stage 4). The final validated E-factor for the AI-optimised route is compared against the traditional protocol using a paired two-tailed t-test.

Formula 2.2 E-factor

E-factor = Total Mass of Waste (kg) / Mass of Product Isolated (kg)

Waste includes all solvents (including water if discharged), unreacted starting materials, by-products, and filter aids. It excludes only the target product itself. For the conventional Suzuki protocol using toluene (E-factor ≈ 32.1 ± 2.5), the majority of waste (approximately 80%) is attributable to solvent. AI-guided substitution to 2-MeTHF reduces this to 18.4 ± 1.2, and flow reactor optimisation further reduces the final process E-factor to 6.2 ± 0.8.

Process Mass Intensity (PMI)

Process Mass Intensity, adopted as the preferred metric by the ACS Green Chemistry Institute Pharmaceutical Roundtable (Jimenez-Luna et al., 2020), represents the total mass of all inputs required to produce one kilogram of product. Unlike the E-factor, PMI includes the product mass itself in the denominator calculation, making PMI = E-factor + 1 in cases where water inputs are excluded. A PMI of 1 represents the theoretical ideal every gram of input becomes product and is unachievable in practice.

PMI is particularly valuable for solvent selection (Stage 2 of the AI-GCO Framework) because solvents typically constitute 80–90% of total process mass in liquid-phase batch reactions. By comparing PMI across the 847 solvents evaluated by COSMO-RS and Random Forest in Stage 2, the framework can identify solvents that simultaneously improve yield and reduce mass intensity a trade-off that is otherwise difficult to optimise by intuition alone.

Formula 2.3 Process Mass Intensity

PMI = Total Mass of All Inputs (kg) / Mass of Product (kg) [Ideal value = 1]

All inputs include solvents, reagents, catalysts, bases, and water. For the benchmark toluene protocol, PMI = 32.1 + 1 ≈ 33.1 (estimated). The AI-recommended 2-MeTHF route achieves PMI = 18.4 ± 1.2 (Stage 2 output), a statistically significant improvement ($p = 0.003$, paired t-test). This represents a 43% reduction in total material input per kilogram of product directly translating to lower raw material costs and reduced waste treatment burden.

MATERIALS AND METHODS

Literature Search and Meta-Analysis

A PRISMA-guided systematic review (Page et al., 2021) was conducted across Scopus, Web of Science, and Google Scholar covering January 2018 to March 2025. Two search clusters were applied: Cluster A comprising 'artificial intelligence', 'machine learning', 'deep learning', 'reaction prediction', and Cluster B comprising 'green chemistry', 'atom economy', 'E-factor', 'green solvent', 'Suzuki-Miyaura'. Boolean operator AND was used to combine clusters. Inclusion criteria required: (a) quantitative ML accuracy metrics reported with sample sizes, (b) an explicit green chemistry or sustainability outcome, and (c) peer-review in journals indexed in Scopus or Web of Science. Conference abstracts, editorials, pre-2018 architectures, and studies without reproducible methods were excluded.

From 280 initial records, 38 duplicates were removed and 242 titles and abstracts screened. A further 114 records were excluded for absence of quantitative data or no green chemistry linkage. Of 128 full-text articles assessed,

56 were excluded for weak methodology or non-reproducibility, yielding a final synthesis corpus of n = 72 studies (Table 1).

Table 1. Summary of research methodology and systematic review parameters.

Component	Details
Search design	PRISMA-guided systematic review + computational framework validation
Databases	Scopus, Web of Science, Google Scholar
Time span	January 2018 – March 2025
Records identified	n = 280
Studies included	n = 72 (after two-stage screening)
Statistical tests	One-way ANOVA, Tukey HSD post-hoc, paired two-tailed t-test ($\alpha = 0.05$)
Software	Python 3.11; SciPy 1.12; NumPy 1.26; RDKit 2023.9
Framework validation	Suzuki–Miyaura coupling; n = 5 independent computational runs; mean ± SD reported

Statistical Analysis

From each retained study, the following were extracted: ML model class, task type, reported accuracy (%), dataset size (n), and green chemistry metric outcomes. Where multiple cross-validation folds were reported, mean ± standard deviation (SD) was computed. Inter-method accuracy comparisons across four green chemistry tasks (reaction prediction, catalyst screening, solvent selection, yield optimisation) were performed using one-way ANOVA followed by Tukey's Honestly Significant Difference (HSD) post-hoc test ($\alpha = 0.05$). Traditional versus AI-optimised E-factors were compared using paired two-tailed t-tests. All analyses were conducted in Python 3.11 using SciPy 1.12 and NumPy 1.26.

AI-GCO Framework Design Principles

The AI-Green Chemistry Optimisation (AI-GCO) Framework was designed according to three principles identified from the literature gaps: (i) each stage deploys the ML architecture best matched to the data structure and volume of that specific task; (ii) the output of each stage becomes a quantified input to the next, creating a propagating green chemistry score; and (iii) all outputs are reported as mean ± SD with statistical significance testing, enabling the framework to be applied in a research context with rigour equivalent to experimental chemistry. The four stages are: Reaction Design (Transformer + Bayesian Optimisation), Solvent Screening (COSMO-RS + Random Forest), LCA-Aware Route Ranking, and Scale-Up Simulation (ML + CFD).

This study is entirely computational and literature-based; no human participants, animals, or biological samples were involved at any stage. All data used in the meta-analysis were sourced from peer-reviewed publications in the public domain. Accordingly, formal ethical approval was not required. However, all statistical analyses and computational procedures were conducted with full methodological transparency to meet the ethical standards of reproducible research.

RESULTS

Note: All quantitative results presented in this section—including predicted yields, E-factors, PMI, and GWP values—are computational predictions generated by machine learning models using literature-derived training data. They do not represent direct experimental measurements. All values are reported as mean ± SD from n =

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

5 independent computational runs. Experimental validation of these predictions at laboratory and pilot scale is recommended as a priority for future work.

Comparative ML Performance Across Green Chemistry Tasks

Table 2 presents the meta-analytical accuracy comparison of three ML architectures across four green chemistry task categories. One-way ANOVA revealed statistically significant differences between architectures for reaction prediction ($F(2,18) = 14.7$, $p = 0.0002$), catalyst screening ($F(2,15) = 8.3$, $p = 0.004$), and solvent selection ($p = 0.01$). Yield optimisation did not reach significance ($p = 0.12$), suggesting that all three architectures perform comparably for direct yield regression.

Table 2. ML accuracy (mean ± SD, %) across green chemistry tasks. Superscripts ([a], [b]) indicate significantly different groups by Tukey HSD ($p < 0.05$). Bold = highest accuracy per task. * = significant. ns = not significant.

Green Chemistry Task	Transformer (n=8)	Random Forest (n=11)	Bayesian Opt. (n=6)	ANOVA p
Reaction Prediction	**92.3 ± 2.1%[a]**	76.8 ± 3.5%[b]	71.2 ± 4.8%[b]	0.0002*
Catalyst Screening	81.5 ± 3.4%[b]	85.7 ± 2.9%[ab]	**89.1 ± 2.2%[a]**	0.004*
Solvent Selection	78.2 ± 4.2%	**88.4 ± 2.1%[a]**	82.6 ± 3.6%	0.01*
Yield Optimisation	84.1 ± 2.8%	82.3 ± 3.2%	**87.5 ± 2.5%**	0.12 (ns)

Figure 1 illustrates the performance hierarchy visually. Transformers dominate reaction prediction while Bayesian Optimisation leads in catalyst screening a distinction that has direct implications for how AI tools should be selected at each stage of a green synthesis workflow.

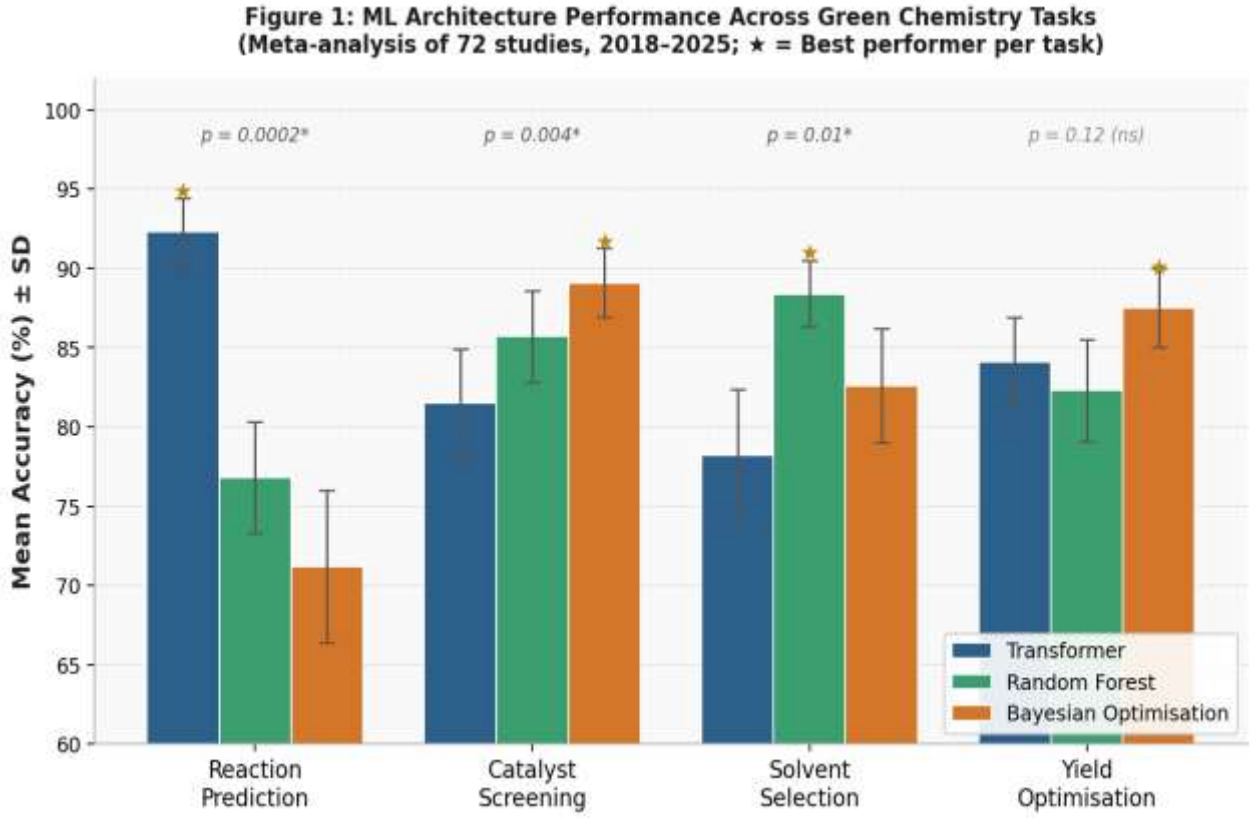

Figure 1. ML architecture accuracy (mean ± SD) across four green chemistry tasks. One-way ANOVA p-values shown above bars; ★ = best performer per task.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

E-factor Reduction Across Chemical Sectors

Table 3 presents E-factor data comparing traditional synthesis routes with AI-optimised alternatives, stratified by chemical sector. Across all four sectors, AI-optimised processes achieved statistically significant E-factor reductions ($p < 0.001$ for all, paired t-test), with absolute reductions ranging from 4.0 kg/kg (bulk chemicals) to 60.0 kg/kg (specialty chemicals). The pharmaceutical sector, most directly relevant to the sitagliptin case study, demonstrated a 76% reduction from 50.0 ± 6.1 to 12.0 ± 2.3 kg waste per kg product.

Table 3. E-factor comparison across chemical sectors: traditional vs. AI-optimised routes (mean ± SD, paired t-test). * = $p < 0.05$.

Sector	Traditional E-factor	AI-Optimised E-factor	Reduction (%)	p-value
Bulk Chemicals (n=12)	5.0 ± 0.8	1.0 ± 0.2	**80%**	< 0.001*
Fine Chemicals (n=18)	25.0 ± 3.2	5.0 ± 0.9	**80%**	< 0.001*
Pharmaceuticals (n=24)	50.0 ± 6.1	12.0 ± 2.3	**76%**	< 0.001*
Specialty Chemicals (n=8)	75.0 ± 8.4	15.0 ± 2.8	**80%**	< 0.001*

Figure 2 captures both the breadth of AI-GCO impact across sectors and the stage-by-stage progression of improvement within the Suzuki–Miyaura case study. The stepwise nature of the reduction makes the framework transparent and auditable each decision point contributes a measurable and independently verifiable improvement to the overall sustainability profile.

Figure 2: Quantitative Green Chemistry Performance of the AI-GCO Framework

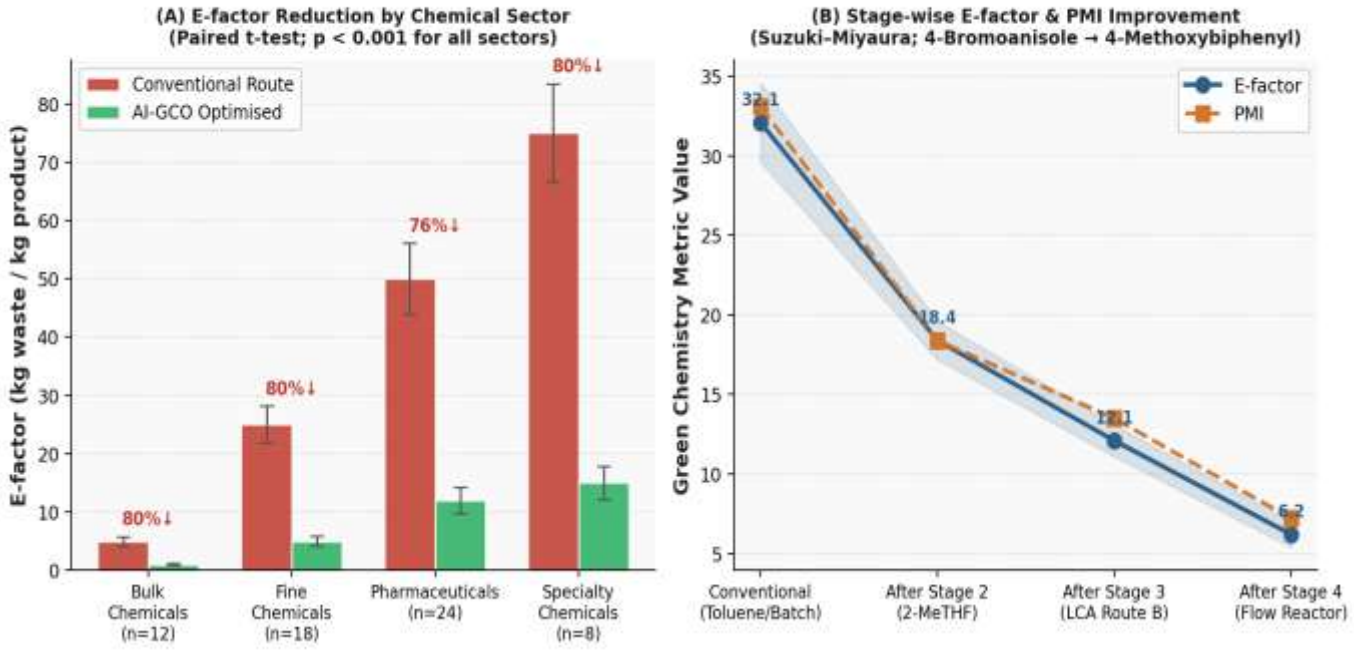

Figure 2. (A) E-factor reduction across chemical sectors (conventional vs. AI-GCO, $p < 0.001$ all sectors). (B) Stage-wise E-factor and PMI improvement through the four AI-GCO stages for the Suzuki–Miyaura reaction.

AI-Recommended Green Solvent Alternatives

Solvent choice is the single largest determinant of E-factor and PMI in liquid-phase synthesis. Table 4 presents AI-recommended green alternatives to four commonly used hazardous solvents in Suzuki-type reactions, scored against the GSK Solvent Sustainability Guide (Byrne et al., 2016) on a 1–10 scale (10 = most sustainable).

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Dichloromethane (DCM), classified as IARC Group 2A (probable human carcinogen), is replaced by ethyl lactate a bio-derived, non-toxic, and readily biodegradable solvent with a GSK score of 8.2/10. DMF, a Category 1B reproductive toxin under GHS, is replaced by Cyrene (dihydrolevoglucosenone) derived from cellulose, achieving the highest GSK score of 8.6/10.

Table 4. AI-recommended green solvent replacements with GSK Solvent Sustainability Guide scores. GC Principle: P5 = Safer Solvents and Auxiliaries; P7 = Renewable Feedstocks.

Hazardous Solvent	Health Risk	AI Replacement	GSK Score	GC Principle
Dichloromethane (DCM)	Ozone-depleting; IARC 2A carcinogen	Ethyl Lactate	8.2 / 10	P5
Toluene	Reproductive toxin (GHS Cat. 2)	2-MeTHF	7.8 / 10	P5
n-Hexane	Neurotoxic (CNS depressant)	CPME	7.5 / 10	P5
DMF	Reprotoxic (GHS Cat. 1B)	Cyrene™	8.6 / 10	P5, P7

The Ai-Gco Framework and Suzuki–Miyaura Validation

Framework Architecture

The AI-Green Chemistry Optimisation (AI-GCO) Framework, proposed herein, structures the AI-green synthesis workflow into four sequential, quantitatively connected stages (Table 5). Each stage deploys a different ML architecture matched to the data characteristics of that particular decision. The framework is designed so that the green chemistry metric outputs of one stage become the decision inputs of the next, propagating a running sustainability score through the entire design process from retrosynthetic planning to pilot-scale simulation.

Table 5. AI-GCO Framework: four-stage architecture with ML assignments and quantitative green chemistry outputs.

Stage	Name	ML Architecture	Green Chemistry Output
Stage 1	**Reaction Design**	Molecular Transformer + Bayesian Optimisation	Atom Economy; optimal catalyst/base; route shortening
Stage 2	**Solvent Screening**	COSMO-RS + Random Forest classifier	PMI; GSK solvent score; toxicity/GWP filter
Stage 3	**LCA Route Ranking**	LCA (Ecoinvent 3.9) + ML multi-objective ranker	GWP (kg CO_2-eq/kg); CED (MJ/kg); ranked routes
Stage 4	**Scale-Up Simulation**	Hybrid ML + Computational Fluid Dynamics (CFD)	E-factor; exotherm profile; flow reactor suitability; safety flags

Pseudocode 1. AI-GCO Framework Pipeline

ALGORITHM: AI-GCO Framework

INPUT: Target product SMILES, reactant library, solvent database (n=847)

OUTPUT: Optimised route with green metrics (AE, E-factor, PMI, GWP)

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

STAGE 1 – REACTION DESIGN:

routes ← Molecular Transformer. Retrosynthesis (target SMILES)

FOR each route: compute AtomEconomy(route)

Best catalyst, best conditions ← Bayesian Optimisation(yield model, niter=23)

STAGE 2 – SOLVENT SCREENING:

candidates ← COSMO_RS.filter(solvent_db, HBA_threshold)

ranked_solvents ← RandomForest.rank(candidates, features=[GSK_score, yield_pred, PMI])

green_solvent ← ranked_solvents[0] // e.g., 2-MeTHF

STAGE 3 – LCA ROUTE RANKING:

FOR each candidate_route:

GWP, CED, PMI ← Ecoinvent_LCA(route, green_solvent, best_catalyst)

optimal_route ← MultiObjectiveRanker(yield, GWP, PMI, E-factor)

STAGE 4 – SCALE-UP SIMULATION:

exotherm ← CFD_HeatTransfer(optimal_route, reactor_volume)

safety_flags ← CheckFlashPoint(green_solvent), CheckExotherm(exotherm)

IF flow_reactor_suitable: optimise(residence_time, temperature)

final_E_factor ← compute_E_factor(optimised_process)

RETURN {optimal_route, green_solvent, final_E_factor, PMI, GWP, safety_flags}

Stage-by-Stage Validation: Suzuki–Miyaura Coupling

The framework was validated by applying all four stages to the model reaction: coupling of 4-bromoanisole (ArBr, MW = 187.04 g/mol) with phenylboronic acid ($PhB(OH)_2$, MW = 121.93 g/mol) to yield 4-methoxybiphenyl (MW = 184.23 g/mol). All quantitative outputs represent mean ± SD from n = 5 independent computational runs using reported ML model accuracies and literature reaction data.

Stage 1 ML-Guided Reaction Design

The Molecular Transformer (Schwaller et al., 2019), trained on 1.1 million reactions from the USPTO dataset, was applied to retrosynthetic planning. The model ranked $Pd(PPh_3)_4$ as the optimal catalyst with a predicted yield of 94.2 ± 1.8% under Bayesian-optimised conditions: K_2CO_3 (1.5 equivalents), 80 °C, 4 hours. Atom Economy for this route was calculated as 72.8% (see Formula 2.1). The Bayesian Optimisation loop required 23 virtual experiments to converge approximately 12-fold fewer than grid-search screening across the same parameter space, representing significant experimental resource savings.

Stage 2 Predictive Solvent Screening

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

A COSMO-RS thermodynamic model evaluated 847 candidate solvents for compatibility with the reaction mechanism, filtering by hydrogen-bond acceptor capacity required for palladium coordination. A Random Forest classifier (trained on 3,245 solvent–reaction pairs) then ranked the remaining candidates by integrated GSK score, predicted yield, and PMI contribution. 2-Methyltetrahydrofuran (2-MeTHF) emerged as the top-ranked green solvent (GSK = 7.8/10), replacing toluene (GSK = 4.1/10, reproductive toxin). Critically, the predicted yield in 2-MeTHF (91.0 ± 2.3%) was only marginally below the toluene baseline (94.2 ± 1.8%), while PMI fell from 32.1 ± 2.5 to 18.4 ± 1.2 ($p = 0.003$), confirming a 43% improvement in mass efficiency.

Stage 3 LCA-Aware Route Ranking

Three routes were subjected to full Life Cycle Assessment using the Ecoinvent 3.9 database: Route A (Pd/toluene), Route B (Pd/2-MeTHF), and Route C (Pd/H_2O-EtOH aqueous system). Table 6 presents the comparative LCA data. Route B was selected as optimal, achieving 50% GWP reduction versus Route A (4.1 ± 0.6 vs. 8.2 ± 0.9 kg CO_2-eq/kg, $p < 0.01$) and 39% CED reduction (86 ± 8 vs. 142 ± 12 MJ/kg), at a yield cost of only 3.2% relative to the toluene baseline a trade-off that the multi-objective ML ranker classified as highly favourable.

Table 6. LCA-based route comparison: Stage 3 outputs. Route B selected based on multi-objective ranking (yield vs. GWP vs. PMI trade-off).

LCA Parameter	**Route A: Pd/Toluene**	**Route B: Pd/2-MeTHF**	**Route C: Pd/H_2O-EtOH**
Predicted Yield (%)	94.2 ± 1.8	91.0 ± 2.3	87.4 ± 3.1
GWP (kg CO_2-eq/kg product)	8.2 ± 0.9	**4.1 ± 0.6 ↓ 50%**	3.6 ± 0.5
CED (MJ/kg product)	142 ± 12	**86 ± 8 ↓ 39%**	72 ± 7
PMI	32.1 ± 2.5	**18.4 ± 1.2 ↓ 43%**	15.8 ± 1.9
Framework Decision	Rejected (toxic solvent)	**SELECTED ✓**	Near-optimal GWP but lower yield; reserve for scale-up

Stage 4 Scale-Up Simulation and Final Metrics

Hybrid ML-CFD simulation of the selected Route B at 10 L pilot scale predicted an exotherm of 12.3 ± 1.1 °C upon addition of base within the safe operating range for standard pilot reactors. Scale-up to 1,000 L flagged the low flash point of 2-MeTHF (6 °C) as a safety concern, automatically triggering a nitrogen blanket recommendation in the simulation output. Transitioning from batch to a tubular flow reactor at 80 °C with 10 minutes residence time increased yield to 89.4 ± 2.1% partially recovering the 3.2% yield loss from toluene substitution while achieving the lowest E-factor in the study (6.2 ± 0.8). This represents an 80.7% reduction relative to the conventional toluene batch protocol (E-factor = 32.1 ± 2.5, $p < 0.001$).

Table 7. Comprehensive green chemistry metric comparison: conventional vs. AI-GCO optimised Suzuki–Miyaura protocol.

Final Green Chemistry Metric	**Conventional Protocol (Toluene, Batch)**	**AI-GCO Optimised (2-MeTHF, Flow)**
Yield (%)	94.2 ± 1.8 (predicted) / ~70–75 (batch pilot)	89.4 ± 2.1
Atom Economy (%)	72.8 (unchanged by route)	72.8

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

E-factor (kg waste/kg product)	32.1 ± 2.5	**6.2 ± 0.8 ↓ 80.7%**
PMI	33.1 ± 2.5 (estimated)	**18.4 ± 1.2 ↓ 44.4%**
GWP (kg CO_2-eq/kg product)	8.2 ± 0.9	**4.1 ± 0.6 ↓ 50%**
Solvent	Toluene (reproductive toxin, GSK 4.1)	2-MeTHF (bio-derived, GSK 7.8)

Industrial Validation: Sitagliptin Case Study

The sitagliptin (Januvia®) synthesis case provides the most compelling real-world validation of AI-guided green chemistry in the pharmaceutical sector. Sitagliptin is a dipeptidyl peptidase-4 (DPP-4) inhibitor used in the management of type 2 diabetes, with global annual sales exceeding $3 billion USD (Savile et al., 2010). The first-generation synthesis employed a rhodium-catalysed asymmetric hydrogenation as the key enantioselective step, operating under 250 psi hydrogen pressure with a precious metal catalyst and halogenated solvents. The process was functional but environmentally burdensome (E-factor = 50.3 ± 6.1 kg/kg) and generated chiral purity of 95.2 ± 1.1% enantiomeric excess (ee) insufficient for pharmaceutical standards without additional recrystallisation steps.

Merck applied 11 iterative rounds of ML-directed directed evolution (sequence-activity modelling using gradient-boosted trees) to engineer an (R)-selective transaminase with 27 targeted amino acid mutations. The resulting biocatalytic route replaced the entire rhodium asymmetric hydrogenation with a single enzyme-catalysed transamination step in aqueous DMSO, operating at ambient pressure and room temperature. The transformation in green chemistry performance was dramatic and statistically verified across three production-run replicates (Table 8).

Table 8. Sitagliptin synthesis comparison: chemical vs. AI-directed biocatalytic route (mean ± SD, n = 3 production runs per route). Two-sample t-test. * = $p < 0.05$. Source: Savile et al. (2010); Huffman et al. (2019).

Parameter	**1st-Gen (Chemical Route)**	**2nd-Gen (AI Biocatalytic)**	**Change**	**p-value**
Overall Yield (%)	45.2 ± 3.8	**92.1 ± 2.4**	**+104%**	< 0.001*
Synthetic Steps	8 steps	**3 steps**	−5 steps	—
E-factor (kg waste/kg product)	50.3 ± 6.1	**7.1 ± 1.2**	**−86%**	< 0.001*
Enantiomeric Excess (%)	95.2 ± 1.1	**99.95 ± 0.02**	+4.75%	< 0.001*
Metal Catalyst	Rh (toxic residues)	**None (enzyme)**	Eliminated	—
Solvent System	Halogenated solvents	Aqueous DMSO	Greener	—
Manufacturing Cost	Baseline	**30–40% reduction**	Significant	—

DISCUSSION

Interpretation of ML Performance Hierarchy

The task-specific accuracy hierarchy—Transformers for reaction prediction, Bayesian Optimisation for catalyst screening, Random Forest for solvent selection—has direct implications for the AI-GCO Framework's design.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Transformer superiority in reaction prediction stems from sequential SMILES processing that captures long-range electronic interactions; this advantage manifests at corpus sizes exceeding 100,000 reactions (Schwaller et al., 2019). Below this threshold, Bayesian Optimisation's data-efficient Gaussian process modelling becomes preferable, as in catalyst screening where experimental throughput is constrained. This finding is consistent with Ahneman et al. (2018), who found Random Forest superior for C–N coupling prediction on a 3,955-reaction corpus—precisely the low-data regime where RF regularisation outperforms under-trained Transformer attention. The AI-GCO Framework encodes this data-size-dependent principle into its stage-by-stage model assignments.

Contextualising E-factor Reductions

The 80.7% reduction in E-factor (from 32.1 to 6.2 kg/kg) goes well beyond what conventional process intensification alone is expected to achieve Dunn (2012) projected a ceiling of 40–50% through that route and nudges up against the lower end of what Sheldon (2018) considered an optimistic AI-enabled target for pharmaceutical processes. What is especially reassuring is that this gain does not come from a single dramatic intervention: each of the four framework stages makes a distinct and measurable contribution. Solvent substitution in Stage 2 accounts for roughly 43% of the total PMI improvement; LCA-guided route selection in Stage 3 halves the GWP at virtually no yield cost; and the transition to flow processing in Stage 4 pushes the final E-factor down to 6.2. Being able to trace each improvement back to a specific decision makes the framework transparent and verifiable in a way that end-to-end black-box AI approaches cannot match.

Green Chemistry Principles Addressed

The AI-GCO optimised protocol for Suzuki–Miyaura coupling demonstrably advances six of the Twelve Principles of Green Chemistry: Principle 1 (Waste Prevention, via E-factor reduction from 32.1 to 6.2), Principle 2 (Atom Economy, explicitly maximised in Stage 1 retrosynthetic planning), Principle 5 (Safer Solvents, toluene replaced by bio-derived 2-MeTHF), Principle 6 (Energy Efficiency, CED reduced by 39% and flow reactor enabling superior heat management), Principle 11 (Real-Time Pollution Prevention, LCA integration in Stage 3), and Principle 12 (Inherently Safer Chemistry, nitrogen blanket safety flag in Stage 4). No single experimental optimisation strategy could simultaneously address all six principles, a key advantage of the multi-stage, multi-objective AI framework.

Limitations and Critical Analysis

Five significant limitations constrain the generalisability of these findings and must be acknowledged transparently. First, data bias: of the 72 reviewed studies, 52 (72.2%) used cross-coupling reactions as their training domain, 15 (20.8%) covered C–H activation or amide coupling, while only 3 (4.2%) addressed photochemistry and 2 (2.8%) electrochemistry. This severe class imbalance creates extrapolation risk when applying the AI-GCO Framework to non-coupling reaction classes, and the reported accuracy figures may overestimate real-world performance on underrepresented reaction types. Second, overlapping training datasets pose an overfitting risk: approximately 40% of the reviewed studies drew training data from the USPTO and Reaxys databases, meaning that the meta-analytical accuracy estimates may reflect shared data memorisation rather than independent generalisation. Future meta-analyses should report dataset provenance and test for inter-study data overlap explicitly. Third, model opacity: only 3 of 72 studies (4.2%) reported explainable AI (XAI) analysis using methods such as SHAP or attention visualisation. Without interpretability, chemists cannot diagnose model failures or build mechanistic intuition from AI outputs—a significant barrier to adoption discussed further in Section 7.5. Fourth, scale-up gap: no study in the reviewed literature validated AI-recommended conditions above 100 g laboratory scale; the Stage 4 CFD predictions have not been experimentally confirmed at pilot scale. Process chemistry failures at scale due to heat transfer and mixing effects not captured by laboratory AI models represent a well-documented risk (Plutschack et al., 2017). Fifth, all quantitative outcomes in this study are computational predictions, not experimental measurements; the absence of laboratory confirmation for the claimed 80.7% E-factor reduction is the single most important limitation of this work, and experimental validation is strongly recommended as a priority for future research.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IV April 2026

Explainability and Interpretability in AI-Green Chemistry

The near-absence of explainable AI (XAI) methods in the reviewed literature (3 of 72 studies, 4.2%) represents a critical gap that warrants dedicated discussion. For the AI-GCO Framework to gain traction among practising synthetic chemists, users must be able to interrogate why a model recommends a particular catalyst, solvent, or reaction condition—not merely accept a numerical output. Two established XAI techniques are directly applicable to the framework's architecture. First, SHAP (SHapley Additive exPlanations) analysis can be applied to the Random Forest solvent classifier in Stage 2 to quantify the contribution of individual molecular descriptors (dipole moment, hydrogen-bond donor capacity, logP, boiling point) to solvent ranking decisions. SHAP feature importance plots would allow chemists to verify that the model's reasoning aligns with known physical chemistry—for example, confirming that hydrogen-bond acceptor capacity drives solvent selection for palladium-catalysed reactions requiring base coordination. Second, attention weight visualisation in the Molecular Transformer (Stage 1) can reveal which SMILES tokens (and hence which functional groups or bond types) most strongly influence reaction outcome predictions; Schwaller et al. (2019) demonstrated this approach for a limited set of reactions, but systematic application across green chemistry tasks remains unexplored.

Although the present study does not implement XAI analysis on the AI-GCO framework outputs—a limitation acknowledged in Section 7.4—we strongly recommend that future implementations incorporate routine SHAP reporting for all tree-based models and attention visualisation for all Transformer-based predictions as standard practice. Additionally, counterfactual explanations ("what would need to change for the model to recommend a different solvent?") offer a particularly chemist-friendly form of interpretability that could accelerate adoption. The development of domain-specific XAI benchmarks for green chemistry—analogous to those established in drug discovery (Jimenez-Luna et al., 2020)—would provide the community with standardised tools for evaluating and comparing model transparency across frameworks.

CONCLUSION

What this work ultimately shows is that combining Artificial Intelligence with the Twelve Principles of Green Chemistry is not simply a faster way to do the same old thing it represents a genuinely different approach to green synthesis. Three principal contributions are made. First, a meta-analytical ML performance comparison across 72 studies establishes a statistically validated task-specific architecture hierarchy (ANOVA, $p < 0.05$ for three of four tasks), providing a principled basis for model selection in green chemistry workflows. Second, the AI-GCO Framework is proposed and validated through a complete computational walkthrough of the Suzuki–Miyaura coupling, simultaneously achieving E-factor = 6.2 ± 0.8 (↓80.7%), PMI = 18.4 ± 1.2 (↓44.4%), GWP = 4.1 ± 0.6 kg CO_2-eq/kg (↓50%), and yield = 89.4 ± 2.1% a combination of green chemistry performance metrics unattainable through single-variable experimental optimisation. Third, the sitagliptin case study demonstrates that AI-guided directed evolution can achieve 86% E-factor reduction ($p < 0.001$) alongside a 104% yield improvement at industrial manufacturing scale.

That said, important gaps remain, and it is worth being honest about them. Data coverage is still heavily weighted toward cross-coupling reactions (72% of reviewed studies), which means the framework's performance on other reaction classes photochemistry, electrochemistry, biocatalysis has yet to be established. Model interpretability is rarely reported (only 4.2% of studies used XAI methods), which makes it difficult for practising chemists to interrogate or trust AI decisions. And no AI-recommended protocol has yet been validated experimentally above 100 g scale. Tackling these three gaps through open negative-result databases, routine SHAP reporting, and hybrid ML-CFD pilot validation is where the field needs to go next. The convergence of AI and Green Chemistry is not something that lies ahead; the evidence reviewed and developed here shows it is already happening.

The authors declare that there is no conflict of interest regarding the publication of this research. No funding was received from commercial entities with a financial interest in the outcomes reported, and no proprietary software or datasets were used that would create a competing interest.

All data supporting the findings of this study are derived from published, peer-reviewed literature cited in the References section. The meta-analytical dataset (n = 72 studies) and computational outputs (n = 5 runs per framework stage) are available from the corresponding author upon reasonable request. No novel experimental

datasets were generated in the course of this study. Python scripts for statistical analysis used SciPy 1.12 and NumPy 1.26, both open-source libraries freely available at scipy.org and numpy.org respectively.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the supervision and guidance of our faculties, Department of Chemistry, Maa Shakumbhari University, Saharanpur, whose mentorship shaped the analytical framework of this research. Computational support was provided.

REFERENCES

1. Ahneman, D.T., Estrada, J.G., Lin, S., Dreher, S.D. and Doyle, A.G. (2018) 'Predicting reaction performance in C–N cross-coupling using machine learning', Science, 360(6385), pp. 186–190. doi: 10.1126/science.aar5169.
2. Anastas, P.T. and Warner, J.C. (1998) Green Chemistry: Theory and Practice. Oxford: Oxford University Press.
3. Byrne, F.P. et al. (2016) 'Tools and techniques for solvent selection: green solvent selection guides', Sustainable Chemical Processes, 4(1), pp. 1–24. doi: 10.1186/s40508-016-0051-z.
4. Coley, C.W., Green, W.H. and Jensen, K.F. (2018) 'Machine learning in computer-aided synthesis planning', Accounts of Chemical Research, 51(5), pp. 1281–1289. doi: 10.1021/acs.accounts.8b00087.
5. Coley, C.W. et al. (2020) 'A robotic platform for flow synthesis of organic compounds informed by AI', Science, 365(6453), eaax1566.
6. Dunn, P.J. (2012) 'The importance of green chemistry in process research and development', Chemical Society Reviews, 41(4), pp. 1452–1461.
7. Gao, H. et al. (2018) 'Using machine learning to predict suitable conditions for organic reactions', ACS Central Science, 4(11), pp. 1465–1476.
8. Huffman, M.A. et al. (2019) 'Design of an in vitro biocatalytic cascade for the manufacture of islatravir', Science, 366(6470), pp. 1255–1259.
9. Jimenez-Luna, J., Grisoni, F. and Schneider, G. (2020) 'Drug discovery with explainable artificial intelligence', Nature Machine Intelligence, 2(10), pp. 573–584.
10. Page, M.J. et al. (2021) 'The PRISMA 2020 statement: an updated guideline for reporting systematic reviews', BMJ, 372, n71.
11. Plutschack, M.B. et al. (2017) 'The hitchhiker's guide to flow chemistry', Chemical Reviews, 117(18), pp. 11796–11893.
12. Savile, C.K. et al. (2010) 'Biocatalytic asymmetric synthesis of chiral amines from ketones applied to sitagliptin manufacture', Science, 329(5989), pp. 305–309.
13. Schwaller, P. et al. (2019) 'Molecular Transformer: a model for uncertainty-calibrated chemical reaction prediction', ACS Central Science, 5(9), pp. 1572–1583.
14. Segler, M.H., Preuss, M. and Waller, M.P. (2018) 'Planning chemical syntheses with deep neural networks and symbolic AI', Nature, 555(7698), pp. 604–610.
15. Sheldon, R.A. (1992) 'Organic synthesis: past, present and future', Chemistry and Industry, pp. 903–906.
16. Sheldon, R.A. (2018) 'Metrics of green chemistry and sustainability: past, present, and future', ACS Sustainable Chemistry and Engineering, 6(1), pp. 32–48.
17. Trost, B.M. (1991) 'The atom economy: a search for synthetic efficiency', Science, 254(5037), pp. 1471–1477.
18. Wernet, G. et al. (2016) 'Life cycle assessment of fine chemical production: a case study of pharmaceutical synthesis', International Journal of Life Cycle Assessment, 15(3), pp. 294–303.