Dynamic large language model representations for multi-objective chemical reaction optimisation

Anonymous Authors[1]

Abstract

Optimising chemical reactions over multiple objectives, such as yield, selectivity, and cost, remains a central challenge in chemical synthesis. Current model-driven approaches rely on fixed molecular representations such as one-hot encodings that discard chemical knowledge, or molecular descriptors that require domain-specific selection workflows and lack common representations across chemically distinct reaction components. Here, we present a multi-objective Bayesian optimisation framework that replaces fixed featurisations with trainable large language model (LLM) embeddings. Textual descriptions of reaction conditions are encoded by a fine-tuned language model jointly optimised with Gaussian process surrogates, producing dense, task-adaptive representations that improve as more experimental data is collected. We benchmark our approach against established descriptor libraries and one-hot encoding baselines across sequential low-data optimisation of nickel- and palladium-catalysed cross-couplings and large-batch 96-well plate high-throughput experimentation regimes, achieving practical convergence thresholds in fewer iterations across all settings. Our framework provides a general, out-of-the-box approach to multi-objective reaction optimisation broadly applicable across reaction types and reaction components, without requiring descriptor selection workflows or domain-specific feature engineering.

1. Introduction

Chemical reaction optimisation is essential to chemical synthesis, and optimising reaction conditions requires navigating vast combinatorial spaces of reaction components such as ligands, catalysts, and bases, while balancing multiple objectives such as yield, selectivity, and cost (Taylor et al.; Düker et al., 2026; Ball et al., 2025; Nippa et al., 2025; Schoepfer et al., 2024). This challenge is particularly acute in pharmaceutical process and preclinical chemistry, where identifying robust conditions for active pharmaceutical ingredient (API) synthesis entails additional environmental, health, and safety considerations (Zhang; Lipton & Barrett, 2006; Zhao et al., 2024; Clarke et al., 2018), and in academic settings where enabling novel transformations often requires extensive exploration of many diverse reaction parameters (Rowsell et al., 2024; Delaney et al., 2023).

Machine learning approaches have proved remarkably effective, with Bayesian optimisation emerging as the leading framework for data-efficient reaction optimisation, having been applied to low-data regimes (Shields et al.; Torres et al.; Braconi & Godineau, 2023; Romer et al., 2024; Cadge et al., 2025b; Zhang et al., 2024; Taylor et al., 2023; Dunlap et al., 2023) and, more recently, to highly parallel high-throughput batched reaction optimisation (Sin et al.; Schlama et al., 2026). Despite these algorithmic advances, how best to represent reaction components for model-driven optimisation remains an open and consequential challenge.

Molecular descriptors have been the dominant featurisation strategy in reaction optimisation and prediction (Figure 1a). Pre-computed descriptor libraries such as Kraken (Gensch et al.) and the COSMO-RS (Moity et al.) database provide readily available steric and electronic properties for established compound classes and have been widely adopted (Dalton et al.; Sin et al.; Schlama et al., 2026; Romer et al., 2024; Cadge et al., 2025b;a; Shved et al., 2024; Christensen et al., 2021). Alternatively, tailored and often laborious feature engineering workflows, involving conformer ensemble analysis, statistical decorrelation, and reaction-specific mechanistic interrogation, can yield highly informative descriptors, though at the cost of computationally expensive density functional theory (DFT) or semi-empirical calculations (Samha et al., 2024; Dotson et al., 2022; Souza et al., 2025; Gandhi et al., 2025; Rinehart et al., 2023; Call et al., 2025; Ocampo et al., 2025; Bannwarth et al., 2019; Kohn et al., 1996).

Despite their widespread use, descriptor-based approaches

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Submitted to the AI for Science workshop (ICML 2026). Do not distribute.

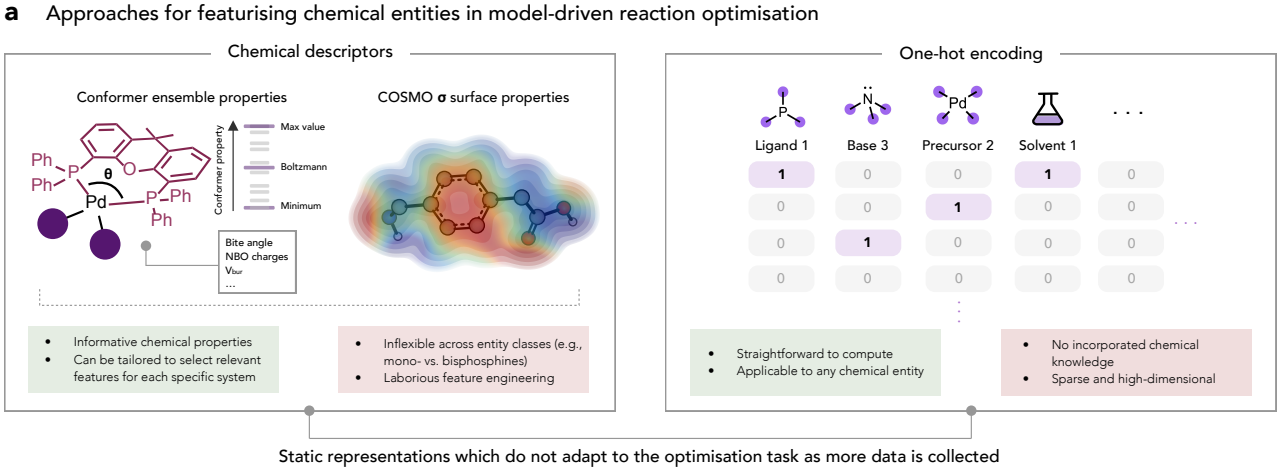

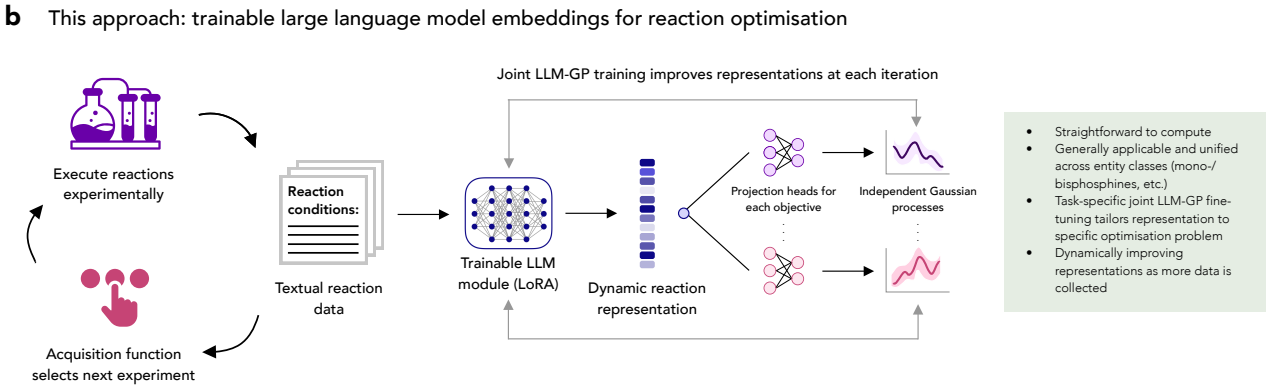

Figure 1. **Featurisation approaches for model-driven reaction optimisation. a,** Conventional approaches for representing chemical entities. Left: molecular descriptors (e.g., Sterimol parameters, buried volume, and bite angle) incorporate physicochemical properties derived from density functional theory (DFT) (Kohn et al., 1996) or extended tight-binding (xTB) (Bannwarth et al., 2019), enabling chemical similarity-aware representations. However, descriptor sets are component-class-specific (e.g., parameters for monodentate ligands are not directly applicable to bidentate ligands), require domain-specific expertise for feature engineering, and can be computationally expensive to generate. Right: one-hot encoding assigns a binary indicator vector to each categorical component (e.g., different ligands, bases, and precursors), yielding sparse, high-dimensional representations that encode no chemical similarity. **b,** This approach: trainable large language model (LLM) embeddings for reaction optimisation. Textual descriptions of reaction conditions are encoded by a LoRA-fine-tuned (Hu et al., 2021) language model into dense, continuous embeddings. These embeddings are passed through objective-specific projection heads into independent Gaussian process (GP) surrogate models, and the entire architecture is trained end-to-end. The architecture is embedded within a multi-objective Bayesian optimisation cycle, with the LLM representations improving at each iteration, organising representations by reaction performance. This yields a unified, task-adaptive representation without requiring descriptor computation or domain-specific feature engineering.

face fundamental limitations. Selecting which descriptors to compute relies on a priori chemical intuition about which molecular properties govern reactivity, a process that is biased by existing mechanistic understanding and must often be repeated for each new reaction system. More critically, descriptors are often not transferable across chemically distinct reaction components. Steric and electronic parameters for monodentate phosphines, for example, are only partially applicable to bidentate ligands, and reactions involving chemically heterogeneous components such as inorganic bases and organometallic additives may lack any meaningful shared descriptor space, making feature engineering increasingly prohibitive as chemical diversity grows.

One-hot encoding offers a simpler alternative that is universally applicable and has shown competitive performance in some settings (Ranković et al., 2024; Shields et al.; Torres et al.). However, it produces sparse, high-dimensional representations that encode no chemical similarity, and its dimensionality scales with the number of reaction components (Figure 1a).

Advances in large language models (LLMs) have shown strong performance across chemical tasks (Ashyrmamatov et al., 2026; Oarga et al., 2026; M. Bran et al., 2024; Zimmermann et al., 2025; Lu & Zhang, 2022; Alampara et al., 2026; Bran et al., 2026; Armstrong et al., 2025) and demon-

strated that their learned embeddings can serve as dense, informative representations for predictive modelling (Ross et al., 2022; Chithrananda et al., 2020; Jablonka et al., 2024). Recent work has further shown that LLM representations can be adapted to specific optimisation tasks through joint fine-tuning with Gaussian process surrogates (Ranković et al., 2025). Crucially, LLMs encode arbitrarily complex chemical information into fixed-length dense representations directly from text, offering a naturally unified featurisation method across any reaction component. These important advances notwithstanding, a general framework for multi-objective reaction optimisation, applicable across chemically heterogeneous reaction components and validated across experimental regimes from low-data to high-throughput settings, has yet to be realised.

Here, we introduce a multi-objective chemical reaction optimisation framework that uses trainable LLM embeddings as a unified reaction representation (Figure 1b). Textual descriptions of reaction conditions are encoded by a low-rank adaptation (LoRA)-fine-tuned language model and passed through reaction objective-specific projection heads into independent Gaussian process surrogates, with the architecture trained end-to-end via joint marginal log-likelihood maximisation. This joint training produces dynamic embeddings that are refined at each optimisation iteration, progressively organising the latent reaction space by reaction performance. We benchmark our approach against established descriptor libraries and one-hot encoding baselines across sequential low-data optimisation of nickel- and palladium-catalysed cross-couplings and large-batch 96-well plate high-throughput experimentation campaigns, demonstrating that our approach achieves practical convergence in fewer iterations across all settings. By replacing fixed, domain-specific featurisations with a dynamic out-of-the-box representation, we provide a general approach to multi-objective reaction optimisation compatible with modern HTE automation platforms, without requiring descriptor computation or feature engineering.

2. Methods

2.1. Metrics for multi-objective optimisation problems

We consider the simultaneous maximisation of M objective functions $y_m \colon \mathcal{X} \to \mathbb{R}$, $m = 1, \ldots, M$ (e.g., reaction yield and selectivity), over a finite chemical design space $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$, where each $\mathbf{x}_i$ is a vector of reaction conditions. The aim is to identify the Pareto optimal set, also referred to as the Pareto front. We say that $\mathbf{x}'$ *dominates* $\mathbf{x}$ (written $\mathbf{x}' \succ \mathbf{x}$) if $y_m(\mathbf{x}') \geq y_m(\mathbf{x})$ for all $m = 1, \ldots, M$ and $y_j(\mathbf{x}') > y_j(\mathbf{x})$ for some $j \in \{1, \ldots, M\}$. The Pareto optimal set is then defined as:

$$\mathcal{P}^* = \left\{ \mathbf{x} \in \mathcal{X} \;\middle|\; \nexists\, \mathbf{x}' \in \mathcal{X} : \mathbf{x}' \succ \mathbf{x} \right\} \tag{1}$$

We seek to approximate $\mathcal{P}^*$ using as few experimental iterations or evaluations as possible with multi-objective Bayesian optimisation. The quality of a candidate Pareto front $\mathcal{P} \subseteq \mathbb{R}^M$, obtained from the current set of observations, is measured by its dominated hypervolume (Figure 2a) relative to a reference point $\mathbf{r} \in \mathbb{R}^M$, chosen such that $r_m < p_m$ for all $\mathbf{p} \in \mathcal{P}$ and $m = 1, \ldots, M$:

$$\mathrm{HV}(\mathcal{P}, \mathbf{r}) = \mathrm{Vol}\Big(\big\{ \mathbf{y} \in \mathbb{R}^M \mid \exists\, \mathbf{p} \in \mathcal{P} : r_m \leq y_m \leq p_m \;\forall\, m \big\} \Big) \tag{2}$$

where $\mathrm{Vol}(\cdot)$ denotes the Lebesgue measure. Hypervolume is a widely used quality indicator for multi-objective optimisation, rewarding both convergence to the true Pareto front and diversity along it (Torres et al.; Zhang et al., 2024; Daulton et al., a;b). Importantly, it is monotonic and Pareto-compliant (Guerreiro et al.; Audet et al., 2021), meaning that if one candidate Pareto front strictly dominates another, it will achieve a strictly higher hypervolume. In this work, we report the normalised hypervolume percentage, defined as $\mathrm{HV}(\mathcal{P}, \mathbf{r}) / \mathrm{HV}(\mathcal{P}^*, \mathbf{r}) \times 100\%$, to evaluate the quality of candidate Pareto fronts $\mathcal{P}$ identified by optimisation algorithms relative to the true Pareto front $\mathcal{P}^*$.

2.2. Representing reaction conditions

We evaluate three featurisation approaches for representing reaction conditions.

One-hot encoding. As a universally applicable and inexpensive baseline that has shown competitive performance in prior work (Ranković et al., 2024; Shields et al.; Torres et al.), each unique categorical variable (e.g., XPhos, dioxane, PhMe) is represented as a binary indicator vector, and all component vectors are concatenated to form the input representation (Figure 1a).

Molecular descriptors. Widely adopted as the primary featurisation strategy in reaction optimisation, molecular descriptors from established descriptor libraries were included as a chemically informative baseline (Dalton et al.; Sin et al.; Schlama et al., 2026; Romer et al., 2024; Cadge et al., 2025b;a; Shved et al., 2024; Christensen et al., 2021). Monophosphine ligands were represented using 190 DFT descriptors from the Kraken library (Gensch et al.), with Principal Component Analysis (PCA) (Maćkiewicz & Ratajczak, 1993) applied to retain components explaining 99% of the variance, reducing dimensionality while preserving information. Solvents were parameterised using four COSMOtherm-derived DFT descriptors from the COSMO-

RS database (Moity et al.). The remaining categorical variables (e.g., bases and precursors), for which comparable descriptor libraries are not readily available, were represented using one-hot encoding.

Large language model representations. Natural language provides a flexible modality for encoding chemically heterogeneous reaction components without requiring component-class-specific featurisation. Each reaction condition $\mathbf{x}_i \in \mathcal{X}$ is represented as a textual prompt describing its components following the template:

```
Reaction condition: ligand: {ligand name} solvent: {solvent name} precursor: {precursor name} base: {base name}
```

A pre-trained language model h_ϕ maps the tokenised prompt to a sequence of embedding vectors $\mathbf{H}_i = h_\phi(\mathbf{x}_i) \in \mathbb{R}^{L \times d_{\text{emb}}}$, where L is the sequence length and d_{emb} is the model's hidden dimension. A pooling operator $\textsc{Pool}: \mathbb{R}^{L \times d_{\text{emb}}} \to \mathbb{R}^{d_{\text{emb}}}$ aggregates the token-level representations into a fixed-dimensional embedding $\mathbf{e}_i = \textsc{Pool}(\mathbf{H}_i) \in \mathbb{R}^{d_{\text{emb}}}$ to produce a sequence-level representation. We adopt pooling strategies based on each model's architecture. For encoder-decoder models, only the encoder stack is used, and pooling is applied over its hidden states. The T5 encoder-decoder models (T5-base (Raffel et al., 2020), T5-small, and T5-chem (Christofidellis et al., 2023)) use mean pooling, which averages hidden states over non-padded tokens. We use last-token pooling for the decoder-only model Qwen2.5 (Yang et al., 2024), which takes the hidden state at the last non-padding position, as causal attention ensures that only this token has attended to the full input sequence (Radford et al., 2018). Mean pooling is also evaluated for Qwen2.5. For the BART-base encoder-decoder model (Lewis et al., 2019), we use mean pooling following the T5 models, and additionally evaluate CLS pooling as BART inherits a `<s>` classification token. All models were accessed via the Hugging Face Transformers library (Wolf et al., 2019). Ablation studies over pre-trained language model architectures and pooling strategies across all benchmark datasets are presented in Appendix Section B.

2.3. Overview of optimisation loop

We initialise our Bayesian optimisation workflows using low-discrepancy Sobol sequences (Burhenne et al.) to ensure diverse coverage of initial points across the search space, which are evaluated to form the initial training data. At each subsequent iteration of the optimisation loop, the surrogate model is fitted to all currently observed data, a batch of candidates is selected by optimising the acquisition function over the remaining design space, and the selected experiments are evaluated and added to the training set.

We use the `qLogNParEGO` acquisition function (Daulton et al., b; Ament et al., 2023; Pilon et al., 2026), a log acquisition function that conditions on noisy baseline observations and applies Chebyshev scalarisation to reduce the multi-objective problem into single-objective subproblems. For batch acquisition, candidates are selected in a greedy sequential fashion, with each selection conditioned on previously chosen pending points.

For surrogate models using fixed input representations (one-hot encoding and molecular descriptors), we used an optimised Gaussian process (GP) configuration adapted from EDBO+ (Torres et al.) and Minerva (Sin et al.) following standard marginal likelihood optimisation. For our approach using an LLM-based deep kernel surrogate with learned input representations, the GP hyperparameters and language model parameters are optimised jointly as described in Section 2.4. All experiments are repeated over 20 random seeds to report statistical variation.

2.4. LLM-based deep kernel Gaussian process surrogate

The LLM first maps tokenised reaction condition prompts to pooled embeddings $\mathbf{e}_i \in \mathbb{R}^{d\text{emb}}$. For each objective $m = 1, \ldots, M$, a trainable projection head g_{ϕ_m} transforms the language model embedding $\mathbf{e}_i$ to an objective-specific representation $\mathbf{z}_i^{(m)} = g_{\phi_m}(\mathbf{e}_i)$. A Gaussian process (GP) then models the mapping from each projected representation to its corresponding objective value:

$$\begin{aligned} y_m(\mathbf{x}_i) &= f_m(\mathbf{z}_i^{(m)}) + \epsilon_m, \\ f_m &\sim \mathcal{GP}(\mu_m, k_{\theta_m}), \\ \epsilon_m &\sim \mathcal{N}(0, \sigma_m^2) \end{aligned} \tag{3}$$

where f_m is the latent function for objective m, modelled as a Gaussian process with constant mean function μ_m and kernel function k_{θ_m}, and ϵ_m is Gaussian observation noise with variance σ_m^2. We use the Matérn-5/2 kernel:

$$k_{\theta_m}(\mathbf{z}, \mathbf{z}') = \sigma_{f,m}^2 \left(1 + \frac{\sqrt{5}\rho}{\ell_m} + \frac{5\rho^2}{3\ell_m^2}\right) \exp\left(-\frac{\sqrt{5}\rho}{\ell_m}\right), \quad \text{where } \rho = \|\mathbf{z} - \mathbf{z}'\|_2 \tag{4}$$

where $\sigma_{f,m}^2$ is the signal variance, and ℓ_m is the length-scale. Together with the constant mean function μ_m and noise variance σ_m^2, these constitute the GP hyperparameters $\theta_m = \{\mu_m, \sigma_{f,m}^2, \ell_m, \sigma_m^2\}$. The kernel matrix for objective m is given by $\mathbf{K}_m = k_{\theta_m}(\mathbf{Z}^{(m)}, \mathbf{Z}^{(m)}) + \sigma_m^2 \mathbf{I}$, where $\mathbf{Z}^{(m)}$ collects the projected embeddings of all observed reaction conditions. The M GPs are queried jointly by the

multi-objective acquisition function. The language model h_ϕ, projection heads $\{g_{\phi_m}\}_{m=1}^M$, and GP hyperparameters $\{\theta_m\}_{m=1}^M$ are trained end-to-end by minimising the sum of negative marginal log-likelihoods:

$$\mathcal{L} = -\sum_{m=1}^{M} \log p(\mathbf{y}_m \mid \mathbf{Z}^{(m)}, \theta_m), \tag{5}$$

where the marginal log-likelihood for objective m is:

$$\log p(\mathbf{y}_m \mid \mathbf{Z}^{(m)}, \theta_m) = -\frac{1}{2}\Big[\tilde{\mathbf{y}}_m^\top \mathbf{K}_m^{-1}\tilde{\mathbf{y}}_m + \log|\mathbf{K}_m| + n\log 2\pi\Big] \tag{6}$$

where n is the number of observations and $\tilde{\mathbf{y}}_m = \mathbf{y}_m - \mu_m\mathbf{1}$ denotes the targets centred by the constant mean function. GP targets are standardised to zero mean and unit variance for training. Gradients from all M objectives are backpropagated through the kernel and projection heads to the shared LLM parameters ϕ, where each projection head g_{ϕ_m} receives gradients only from its corresponding objective. This enables the LLM to learn representations that are jointly optimised for GP predictive performance across all objectives, while each projection head specialises for its target objective.

2.5. Parameter-efficient fine-tuning of language models with LoRA

The language model embeddings of reaction conditions evolve across BO iterations as the LLM is fine-tuned on accumulating observations. To do so efficiently, we apply parameter-efficient fine-tuning (PEFT) with Low-Rank Adaptation (LoRA) (Hu et al., 2021). Rather than updating all of the pre-trained language model's parameters, we use LoRA to update only a small subset. For a pre-trained weight matrix $\mathbf{W}_0 \in \mathbb{R}^{d\times k}$, LoRA injects trainable low-rank decompositions such that the weight update takes the form:

$$\begin{aligned}\mathbf{W} &= \mathbf{W}_0 + \Delta\mathbf{W},\\ \Delta\mathbf{W} &= \frac{\alpha}{r}\mathbf{B}\mathbf{A}, \quad \text{with } \mathbf{B}\in\mathbb{R}^{d\times r}, \mathbf{A}\in\mathbb{R}^{r\times k}\end{aligned} \tag{7}$$

where $r \ll \min(d,k)$ is the rank and α is a scaling hyperparameter. We use LoRA to target a subset of linear projection matrices within the pre-trained language models, using rank $r = 4$ and $\alpha = 16$.

3. Results

We benchmarked our approach (MO-LLM) against established featurisation methods in reaction optimisation: DFT descriptors, computed from the Kraken (Gensch et al.) and COSMO-RS (Moity et al.) descriptor libraries, and one-hot encoding baselines (see Section 2.2 for more details). All reaction representations were embedded within the same multi-objective Bayesian loop and acquisition function (see Section 2.3). For each benchmark dataset, all reaction representation methods were allocated an identical experimental budget, determined based on the number of optimisation iterations required for at least one method to reach practical convergence (≥95% of the maximum hypervolume) on average across 20 random seeds (Figure 2a). We report the mean number of iterations required to reach 90% and 95% hypervolume thresholds, with standard error across seeds, as these thresholds represent practically relevant levels of convergence to the true Pareto front (see Section 2.1 for a more detailed explanation of hypervolume metrics). We additionally report the percentage of optimisation runs (initialised with different random seeds) reaching each hypervolume threshold within the allocated experimental budget as a measure of robustness; these results are provided in Appendix Section A. Among the pre-trained language models and pooling strategies evaluated (T5-base (Raffel et al., 2020), T5-small, T5-chem (Lu & Zhang, 2022), BART-base (Lewis et al., 2019), and Qwen2.5-0.5B (Yang et al., 2024)), T5-base with mean pooling consistently yielded the best optimisation performance and is used in MO-LLM throughout this study. Full ablation studies across model architectures and pooling strategies are provided in Appendix Section B.

3.1. Sequential low-data reaction optimisation

We first evaluated our approach (MO-LLM) in sequential low-data optimisation regimes, where minimising the number of experiments required to identify high-performing reaction conditions is the primary goal (Figure 2b). We used two open-source multi-objective reaction datasets, each comprising experimental yield and selectivity as the optimisation objectives and presenting large combinatorial search spaces with diverse categorical reaction components: a nickel-catalysed Suzuki coupling (384 experiments from a space of 88,000 combinations, varying ligand, Ni-precursor, base, solvent, co-solvent, and temperature) (Sin et al.), and a palladium-catalysed Suzuki coupling (404 experiments from 59,040 combinations, varying ligand, Pd-precursor, base, solvent, and co-solvent) (Schlama et al., 2026). Each campaign was initialised with 5 experiments selected using diversity-guided Sobol sampling (see Section 2.3), after which the Bayesian optimisation loop selected one candidate per iteration. On the nickel-catalysed Suzuki coupling, MO-LLM reached the 90% hypervolume threshold in approximately 13 iterations on average, roughly half the number required by both the DFT descriptor-based (∼26 iterations) and one-hot encoding (∼24 iterations) baselines (Figure 2b). At the more stringent 95% threshold, the gap narrowed, though MO-LLM retained a consistent advantage,

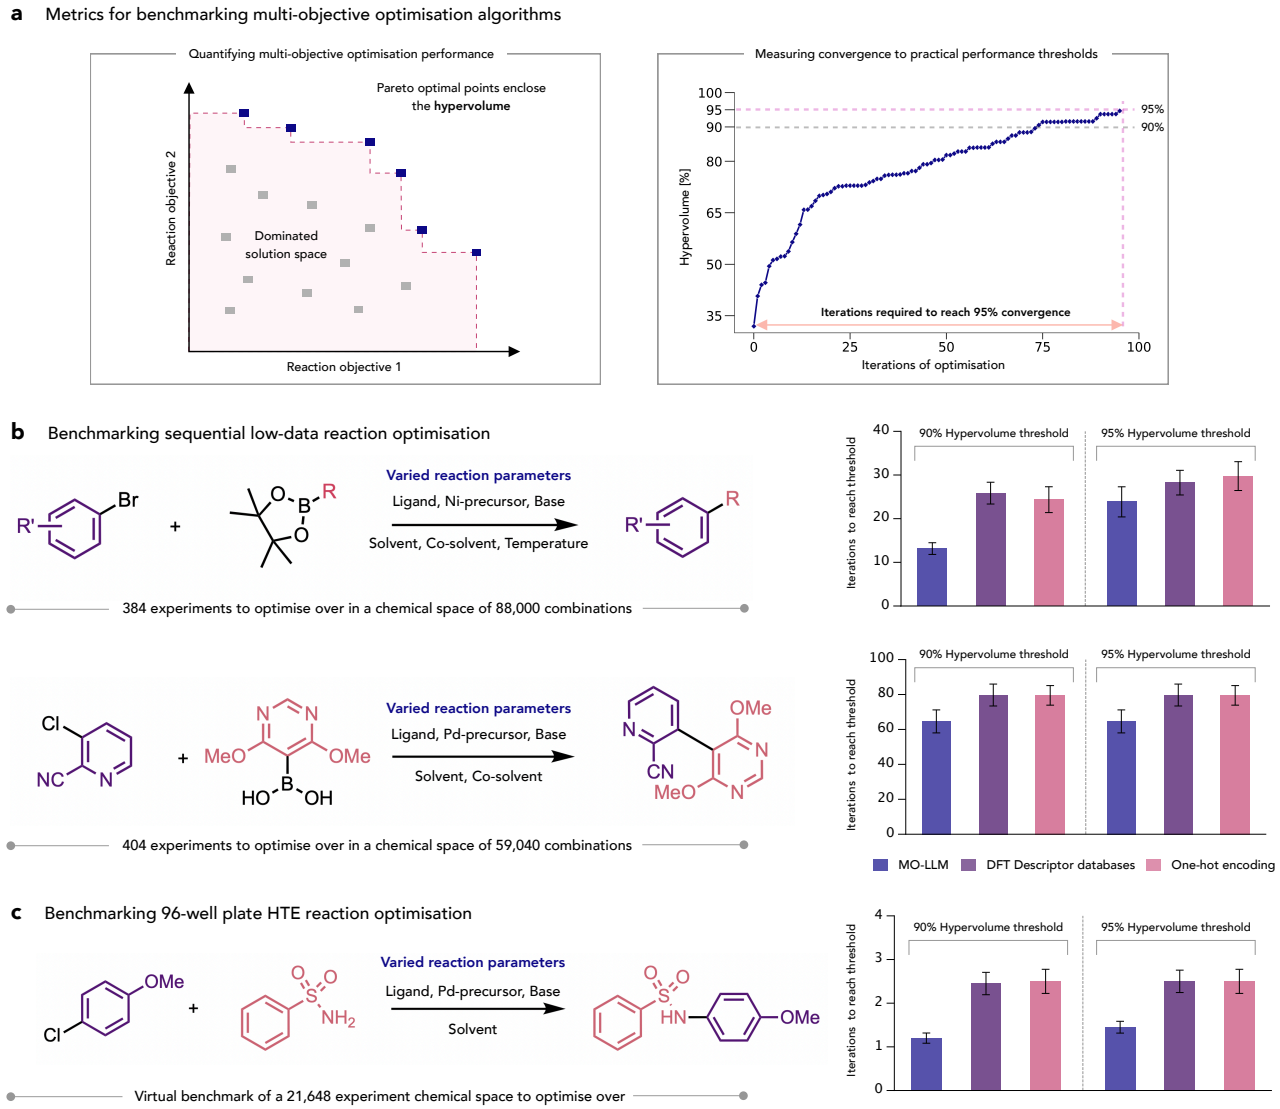

Figure 2. **Benchmarking multi-objective reaction optimisation across experimental regimes. a,** Metrics for evaluating multi-objective optimisation. Left: the hypervolume indicator quantifies the volume of objective space dominated by the current Pareto-optimal set, providing a scalar measure of multi-objective performance. Right: convergence is assessed by tracking the normalised hypervolume over optimisation iterations; the number of iterations required to reach practical performance thresholds (90% and 95% of maximum hypervolume) serves as a comparative metric across featurisation methods. Additionally, we report the percentage optimisation runs (initialised with different random seeds) reaching each threshold as a measure of robustness in Appendix Section A. **b,** Sequential low-data reaction optimisation benchmarks. Two multi-objective case studies from published reaction datasets are evaluated: a nickel-catalysed Suzuki coupling (Sin et al.) (384 experiments from a space of 88,000 combinations) and a palladium-catalysed Suzuki coupling (Schlama et al., 2026) (404 experiments from a space of 59,040 combinations). Iterations required to reach 90% and 95% hypervolume thresholds are compared for our approach (MO-LLM), DFT descriptor libraries (Kraken (Gensch et al.), COSMO-RS (Moity et al.)), and one-hot encoding, repeated across 20 random seeds (see Section 2.2 for more details). MO-LLM consistently reaches both thresholds in fewer iterations across both case studies. **c,** 96-well plate HTE reaction optimisation benchmark. A palladium-catalysed sulfonamide coupling (Schlama et al., 2026) (virtual benchmark of 21,648 experiments) from a published dataset is optimised in large parallel batches of 96 experiments. MO-LLM achieves convergence in fewer plate iterations than descriptor-based and one-hot baselines.

converging in approximately 24 iterations compared to ∼28 for DFT descriptors and ∼30 for one-hot encoding. On the palladium-catalysed Suzuki coupling benchmark, MO-LLM achieved both the 90% and 95% thresholds in approximately 65 iterations on average, while the descriptor-based and one-hot encoding baselines required approximately 80 iterations each to reach the same performance levels (Figure 2b). Across both case studies, MO-LLM consistently reached practical convergence thresholds in fewer experiments than either baseline, improving sample efficiency. MO-LLM also

achieved a consistently higher proportion of successful optimisation runs reaching both convergence thresholds within the allocated experimental budget, indicating more robust convergence across different random seeds (see Appendix Section A).

3.2. Highly parallel 96-well HTE reaction optimisation

Modern high-throughput experimentation (HTE) platforms have combined parallel screening with data-driven optimisation (Sin et al.; Porte et al., 2026), enabling the exploration of large reaction spaces in compressed experimental timescales. We next sought to evaluate MO-LLM in this batched optimisation setting, where reducing the number of HTE plate iterations directly translates to savings in experimental time. We used an open-source palladium-catalysed sulfonamide coupling dataset (21,648 virtual experiments, varying ligand, Pd-precursor, base, and solvent) (Schlama et al., 2026) as a benchmark, optimising yield and selectivity in parallel batches of 96 experiments to simulate 96-well plate HTE campaigns (Figure 2c). Results for 24- and 48-well plate batch sizes are provided in Appendix Section A. MO-LLM reached practical convergence thresholds in approximately 1.2 plate optimisation iterations on average, whereas both the DFT descriptor-based and one-hot encoding baselines required approximately double the number of plate optimisation iterations ($\sim$2.4) to achieve the same thresholds (Figure 2c). MO-LLM also achieved a higher proportion of successful optimisation runs reaching both convergence thresholds, with detailed results provided in Appendix Section A. Given that each 96-well plate campaign typically requires approximately one week of experimental and analytical time, this reduction translates directly into meaningful savings in time and resources. These results demonstrate that our approach offers consistent advantages across both sequential low-data regimes, where sample efficiency is crucial, and parallel large-batch settings increasingly adopted in pharmaceutical process development and academic high-throughput screening (Götz et al., 2023; 2025; Mahjour et al., 2023).

Software and Data

All data and custom code developed for this work is implemented in Python and will be made publicly available in an open-source GitHub repository under the Apache 2.0 licence upon acceptance of this manuscript.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Alampara, N., Aneesh, A., Ríos-García, M., Mirza, A., Schilling-Wilhelmi, M., Aghajani, A. A., Sun, M., Prastalo, G., and Jablonka, K. M. General-purpose models for the chemical sciences: Llms and beyond. *Chemical Reviews*, 126(4):2484–2549, February 2026. ISSN 1520-6890. doi: 10.1021/acs.chemrev.5c00583. URL http://dx.doi.org/10.1021/acs.chemrev.5c00583.

Ament, S., Daulton, S., Eriksson, D., Balandat, M., and Bakshy, E. Unexpected improvements to expected improvement for bayesian optimization, 2023. URL https://arxiv.org/abs/2310.20708.

Armstrong, D., Jončev, Z., Bran, A. M., and Schwaller, P. Synthstrategy: Extracting and formalizing latent strategic insights from llms in organic chemistry, 2025. URL https://arxiv.org/abs/2512.01507.

Ashyrmamatov, I., Gwak, S. J., Jin, S.-Y., Jun, I., Ucak, U. V., Lee, J.-Y., and Lee, J. A survey on large language models in biology and chemistry. *Experimental amp; Molecular Medicine*, April 2026. ISSN 2092-6413. doi: 10.1038/s12276-025-01583-1. URL http://dx.doi.org/10.1038/s12276-025-01583-1.

Audet, C., Bigeon, J., Cartier, D., Le Digabel, S., and Salomon, L. Performance indicators in multiobjective optimization. *European Journal of Operational Research*, 292(2):397–422, 2021. ISSN 0377-2217. doi: 10.1016/j.ejor.2020.11.016. URL http://dx.doi.org/10.1016/j.ejor.2020.11.016.

Ball, M., Horvath, D., Kogej, T., Kabeshov, M., and Varnek, A. Predicting reaction conditions: a data-driven perspective. *Chemical Science*, 16(38):17523–17541, 2025. ISSN 2041-6539. doi: 10.1039/d5sc03045e. URL http://dx.doi.org/10.1039/D5SC03045E.

Bannwarth, C., Ehlert, S., and Grimme, S. Gfn2-xtb—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of Chemical Theory and Computation*, 15(3):1652–1671, February 2019. ISSN 1549-9626. doi: 10.1021/acs.jctc.8b01176. URL http://dx.doi.org/10.1021/acs.jctc.8b01176.

Braconi, E. and Godineau, E. Bayesian optimization as a sustainable strategy for early-stage process development? a case study of cu-catalyzed c–n coupling of sterically hindered pyrazines. *ACS Sustainable Chemistry amp; Engineering*, 11(28):10545–10554, 2023. ISSN 2168-0485. doi: 10.1021/acssuschemeng.3c02455. URL http://dx.doi.org/10.1021/acssuschemeng.3c02455.

Bran, A. M., Neukomm, T. A., Armstrong, D., Jončev, Z., and Schwaller, P. Chemical reasoning in llms unlocks strategy-aware synthesis planning and reaction mechanism elucidation. *Matter*, pp. 102812, April 2026. ISSN 2590-2385. doi: 10.1016/j.matt.2026.102812. URL http://dx.doi.org/10.1016/j.matt.2026.102812.

Burhenne, S., Jacob, D., and Henze, G. P. Sampling based on Sobol' sequences for Monte Carlo techniques applied to building simulations. pp. 1816–1823.

Cadge, J. A., Hart, S. D., Walroth, R. C., Mack, K. A., and Sigman, M. S. Bisphosphine ligand conformer selection to enhance descriptor database representation: improving statistical modelling outcomes. *Chemical Science*, 16 (43):20473–20485, 2025a. ISSN 2041-6539. doi: 10.1039/d5sc04691b. URL http://dx.doi.org/10.1039/D5SC04691B.

Cadge, J. A., Lozano, C., Merriman, M. T., Oblad, P., Sigman, M. S., and Reisman, S. E. A data science-guided approach for the development of nickel-catalyzed homodiels–alder reactions. *Journal of the American Chemical Society*, 147(34):31175–31186, August 2025b. ISSN 1520-5126. doi: 10.1021/jacs.5c09948. URL http://dx.doi.org/10.1021/jacs.5c09948.

Call, A., Palone, A., Liles, J. P., Romer, N. P., Read, J. A., Luis, J. M., Sigman, M. S., Bietti, M., and Costas, M. Understanding catalytic enantioselective c–h bond oxidation at nonactivated methylenes through predictive statistical modeling analysis. *ACS Catalysis*, 15 (3):2110–2123, January 2025. ISSN 2155-5435. doi: 10.1021/acscatal.4c05659. URL http://dx.doi.org/10.1021/acscatal.4c05659.

Chithrananda, S., Grand, G., and Ramsundar, B. Chemberta: Large-scale self-supervised pretraining for molecular property prediction, 2020. URL https://arxiv.org/abs/2010.09885.

Christensen, M., Yunker, L. P. E., Adedeji, F., Häse, F., Roch, L. M., Gensch, T., dos Passos Gomes, G., Zepel, T., Sigman, M. S., Aspuru-Guzik, A., and Hein, J. E. Data-science driven autonomous process optimization. *Communications Chemistry*, 4(1), August 2021. ISSN 2399-3669. doi: 10.1038/s42004-021-00550-x. URL http://dx.doi.org/10.1038/s42004-021-00550-x.

Christofidellis, D., Giannone, G., Born, J., Winther, O., Laino, T., and Manica, M. Unifying molecular and textual representations via multi-task language modelling. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 6140–6157. PMLR, 2023. URL https://proceedings.mlr.press/v202/christofidellis23a.html.

Clarke, C. J., Tu, W.-C., Levers, O., Bröhl, A., and Hallett, J. P. Green and sustainable solvents in chemical processes. *Chemical Reviews*, 118(2):747–800, January 2018. ISSN 1520-6890. doi: 10.1021/acs.chemrev.7b00571. URL http://dx.doi.org/10.1021/acs.chemrev.7b00571.

Dalton, D. M., Walroth, R. C., Rouget-Virbel, C., Mack, K. A., and Toste, F. D. Utopia Point Bayesian Optimization Finds Condition-Dependent Selectivity for N-Methyl Pyrazole Condensation. 146(23):15779–15786. ISSN 0002-7863. doi: 10.1021/jacs.4c01616. URL https://doi.org/10.1021/jacs.4c01616.

Daulton, S., Balandat, M., and Bakshy, E. Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9851–9864. Curran Associates, Inc., a. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/6fec24eac8f18ed793f5eaad3dd7977c-Abstract.html.

Daulton, S., Balandat, M., and Bakshy, E. Parallel Bayesian Optimization of Multiple Noisy Objectives with Expected Hypervolume Improvement, b. URL http://arxiv.org/abs/2105.08195.

Delaney, C. P., Lin, E., Huang, Q., Yu, I. F., Rao, G., Tao, L., Jed, A., Fantasia, S. M., Püntener, K. A., Britt, R. D., and Hartwig, J. F. Cross-coupling by a noncanonical mechanism involving the addition of aryl halide to cu(ii). *Science*, 381(6662):1079–1085, 2023. ISSN 1095-9203. doi: 10.1126/science.adi9226. URL http://dx.doi.org/10.1126/science.adi9226.

Dotson, J. J., van Dijk, L., Timmerman, J. C., Grosslight, S., Walroth, R. C., Gosselin, F., Püntener, K., Mack, K. A., and Sigman, M. S. Data-driven multi-objective optimization tactics for catalytic asymmetric reactions using bisphosphine ligands. *Journal of the American Chemical Society*, 145(1):110–121, December 2022. ISSN 1520-5126. doi: 10.1021/jacs.2c08513. URL http://dx.doi.org/10.1021/jacs.2c08513.

Düker, J., Hebing, L., Leweke, S., Nicholls, R. L., Lübbesmeyer, M., Volpin, G., König, B., and Hillenbrand, J. Data science-assisted workflow for reaction optimization in process chemistry. *Organic Process Research amp; Development*, 30(1):175–188, January 2026. ISSN 1520-586X. doi: 10.1021/acs.oprd.5c00384. URL http://dx.doi.org/10.1021/acs.oprd.5c00384.

Dunlap, J. H., Ethier, J. G., Putnam-Neeb, A. A., Iyer, S., Luo, S.-X. L., Feng, H., Garrido Torres, J. A., Doyle, A. G., Swager, T. M., Vaia, R. A., Mirau, P., Crouse, C. A., and Baldwin, L. A. Continuous flow synthesis of pyridinium salts accelerated by multi-objective bayesian optimization with active learning. *Chemical Science*, 14(30):8061–8069, 2023. ISSN 2041-6539. doi: 10.1039/d3sc01303k. URL http://dx.doi.org/10.1039/D3SC01303K.

Gandhi, S. S., Brown, G. Z., Aikonen, S., Compton, J. S., Neves, P., Martinez Alvarado, J. I., Strambeanu, I. I., Leonard, K. A., and Doyle, A. G. Data science-driven discovery of optimal conditions and a condition-selection model for the chan–lam coupling of primary sulfonamides. *ACS Catalysis*, 15(3):2292–2304, January 2025. ISSN 2155-5435. doi: 10.1021/acscatal.4c07972. URL http://dx.doi.org/10.1021/acscatal.4c07972.

Gensch, T., family=Passos Gomes, given=Gabriel, p. u., Friederich, P., Peters, E., Gaudin, T., Pollice, R., Jorner, K., Nigam, A., Lindner-D'Addario, M., Sigman, M. S., and Aspuru-Guzik, A. A Comprehensive Discovery Platform for Organophosphorus Ligands for Catalysis. 144(3):1205–1217. ISSN 0002-7863. doi: 10.1021/jacs.1c09718. URL https://doi.org/10.1021/jacs.1c09718.

Götz, J., Jackl, M. K., Jindakun, C., Marziale, A. N., André, J., Gosling, D. J., Springer, C., Palmieri, M., Reck, M., Luneau, A., Brocklehurst, C. E., and Bode, J. W. High-throughput synthesis provides data for predicting molecular properties and reaction success. *Science Advances*, 9(43), October 2023. ISSN 2375-2548. doi: 10.1126/sciadv.adj2314. URL http://dx.doi.org/10.1126/sciadv.adj2314.

Götz, J., Richards, E., Stepek, I. A., Takahashi, Y., Huang, Y.-L., Bertschi, L., Rubi, B., and Bode, J. W. Predicting three-component reaction outcomes from 40, 000 miniaturized reactant combinations. *Science Advances*, 11(22), May 2025. ISSN 2375-2548. doi: 10.1126/sciadv.adw6047. URL http://dx.doi.org/10.1126/sciadv.adw6047.

Guerreiro, A. P., Fonseca, C. M., and Paquete, L. The Hypervolume Indicator: Problems and Algorithms. 54(6):1–42. ISSN 0360-0300, 1557-7341. doi: 10.1145/3453474. URL http://arxiv.org/abs/2005.00515.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models, 2021. URL https://arxiv.org/abs/2106.09685.

Jablonka, K. M., Schwaller, P., Ortega-Guerrero, A., and Smit, B. Leveraging large language models for predictive chemistry. *Nature Machine Intelligence*, 6(2):161–169, February 2024. ISSN 2522-5839. doi: 10.1038/s42256-023-00788-1. URL http://dx.doi.org/10.1038/s42256-023-00788-1.

Kohn, W., Becke, A. D., and Parr, R. G. Density functional theory of electronic structure. *The Journal of Physical Chemistry*, 100(31):12974–12980, January 1996. ISSN 1541-5740. doi: 10.1021/jp960669l. URL http://dx.doi.org/10.1021/jp960669l.

Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461, 2019. URL http://arxiv.org/abs/1910.13461.

Lipton, M. F. and Barrett, A. G. M. Introduction: process chemistry. *Chemical Reviews*, 106(7):2581–2582, 2006. ISSN 1520-6890. doi: 10.1021/cr068400d. URL http://dx.doi.org/10.1021/cr068400d.

Lu, J. and Zhang, Y. Unified deep learning model for multitask reaction predictions with explanation. *Journal of Chemical Information and Modeling*, 62(6):1376–1387, March 2022. ISSN 1549-960X. doi: 10.1021/acs.jcim.1c01467. URL http://dx.doi.org/10.1021/acs.jcim.1c01467.

M. Bran, A., Cox, S., Schilter, O., Baldassari, C., White, A. D., and Schwaller, P. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5):525–535, May 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00832-8. URL http://dx.doi.org/10.1038/s42256-024-00832-8.

Mahjour, B., Zhang, R., Shen, Y., McGrath, A., Zhao, R., Mohamed, O. G., Lin, Y., Zhang, Z., Douthwaite, J. L., Tripathi, A., and Cernak, T. Rapid planning and analysis of high-throughput experiment arrays for reaction discovery. *Nature Communications*, 14(1), 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-39531-0. URL http://dx.doi.org/10.1038/s41467-023-39531-0.

Maćkiewicz, A. and Ratajczak, W. Principal components analysis (pca). *Computers amp; Geosciences*, 19(3):303–342, March 1993. ISSN 0098-3004. doi: 10.1016/0098-3004(93)90090-r. URL http://dx.doi.org/10.1016/0098-3004(93)90090-R.

Moity, L., Durand, M., Benazzouz, A., Pierlot, C., Molinier, V., and Aubry, J.-M. Panorama of sustainable

solvents using the COSMO-RS approach. 14(4):1132–1145. ISSN 1463-9270. doi: 10.1039/C2GC16515E. URL https://pubs.rsc.org/en/content/articlelanding/2012/gc/c2gc16515e.

Nippa, D. F., Boddy, A. J., Atz, K., Grether, U., Binch, H., and Martin, R. E. Accelerating compound synthesis in drug discovery: the role of digitalisation and automation. *RSC Medicinal Chemistry*, 16(12):5753–5764, 2025. ISSN 2632-8682. doi: 10.1039/d5md00672d. URL http://dx.doi.org/10.1039/D5MD00672D.

Oarga, A., Hart, M., Bran, A. M., Lederbauer, M., and Schwaller, P. Scientific knowledge graph and ontology generation using open large language models. *Digital Discovery*, 5(3):1269–1279, 2026. ISSN 2635-098X. doi: 10.1039/d5dd00275c. URL http://dx.doi.org/10.1039/D5DD00275C.

Ocampo, B. E., Altundas, B., Bock, M. J., Feiz, S., and Denmark, S. E. Data-driven prediction of enantioselectivity for the sharpless asymmetric dihydroxylation: Model development and experimental validation. *ACS Central Science*, 11(9):1640–1650, 2025. ISSN 2374-7951. doi: 10.1021/acscentsci.5c00900. URL http://dx.doi.org/10.1021/acscentsci.5c00900.

Pilon, S., Savino, E., Bayley, O. M., Vanzella, M., Claros, M., Siasiaridis, P., Liu, J., Lukas, F., Damian, M., Tseliou, V., Intini, N., Slattery, A., SanJosé-Orduna, J., den Hartog, T., Peters, R. A. H., Gargano, A. F. G., Mutti, F. G., and Noël, T. A flexible and affordable self-driving laboratory for automated reaction optimization. *Nature Synthesis*, April 2026. ISSN 2731-0582. doi: 10.1038/s44160-026-01053-0. URL http://dx.doi.org/10.1038/s44160-026-01053-0.

Porte, V., Hepp, L., Kollmus, P., Lu, S., Serrano, E., Blanco, D., and Santagostino, M. Frugal sampling strategies for navigating complex reaction spaces. *Organic Process Research amp; Development*, April 2026. ISSN 1520-586X. doi: 10.1021/acs.oprd.6c00027. URL http://dx.doi.org/10.1021/acs.oprd.6c00027.

Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL http://jmlr.org/papers/v21/20-074.html.

Ranković, B., Griffiths, R.-R., Moss, H. B., and Schwaller, P. Bayesian optimisation for additive screening and yield improvements – beyond one-hot encoding. *Digital Discovery*, 3(4):654–666, 2024. ISSN 2635-098X. doi: 10.1039/d3dd00096f. URL http://dx.doi.org/10.1039/D3DD00096F.

Ranković, B., Griffiths, R.-R., and Schwaller, P. Large language models as uncertainty-calibrated optimizers for experimental discovery, 2025. URL https://arxiv.org/abs/2504.06265.

Rinehart, N. I., Saunthwal, R. K., Wellauer, J., Zahrt, A. F., Schlemper, L., Shved, A. S., Bigler, R., Fantasia, S., and Denmark, S. E. A machine-learning tool to predict substrate-adaptive conditions for pd-catalyzed c–n couplings. *Science*, 381(6661):965–972, 2023. ISSN 1095-9203. doi: 10.1126/science.adg2114. URL http://dx.doi.org/10.1126/science.adg2114.

Romer, N. P., Min, D. S., Wang, J. Y., Walroth, R. C., Mack, K. A., Sirois, L. E., Gosselin, F., Zell, D., Doyle, A. G., and Sigman, M. S. Data science guided multiobjective optimization of a stereoconvergent nickel-catalyzed reduction of enol tosylates to access trisubstituted alkenes. *ACS Catalysis*, 14(7):4699–4708, March 2024. ISSN 2155-5435. doi: 10.1021/acscatal.4c00650. URL http://dx.doi.org/10.1021/acscatal.4c00650.

Ross, J., Belgodere, B., Chenthamarakshan, V., Padhi, I., Mroueh, Y., and Das, P. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, December 2022. ISSN 2522-5839. doi: 10.1038/s42256-022-00580-7. URL http://dx.doi.org/10.1038/s42256-022-00580-7.

Rowsell, B. J. S., O'Brien, H. M., Athavan, G., Daley-Dee, P. R., Krieger, J., Richards, E., Heaton, K., Fairlamb, I. J. S., and Bedford, R. B. The iron-catalysed suzuki coupling of aryl chlorides. *Nature Catalysis*, 7 (11):1186–1198, October 2024. ISSN 2520-1158. doi: 10.1038/s41929-024-01234-0. URL http://dx.doi.org/10.1038/s41929-024-01234-0.

Samha, M. H., Karas, L. J., Vogt, D. B., Odogwu, E. C., Elward, J., Crawford, J. M., Steves, J. E., and Sigman, M. S. Predicting success in cu-catalyzed c–n coupling reactions using data science. *Science Advances*, 10(3), January 2024. ISSN 2375-2548. doi: 10.1126/sciadv.adn3478. URL http://dx.doi.org/10.1126/sciadv.adn3478.

Schlama, R., Sin, J. W., Burwood, R. P., Püntener, K., Bigler, R., and Schwaller, P. Swarm intelligence for chemical reaction optimization. *Chem*, pp. 103035, April 2026. ISSN 2451-9294. doi: 10.1016/j.chempr.2026.103035. URL http://dx.doi.org/10.1016/j.chempr.2026.103035.

Schoepfer, A. A., Weinreich, J., Laplaza, R., Waser, J., and Corminboeuf, C. Cost-informed bayesian reaction optimization. *Digital Discovery*, 3(11):2289–2297, 2024. ISSN 2635-098X. doi: 10.1039/d4dd00225c. URL http://dx.doi.org/10.1039/D4DD00225C.

Shields, B. J., Stevens, J., Li, J., Parasram, M., Damani, F., Alvarado, J. I. M., Janey, J. M., Adams, R. P., and Doyle, A. G. Bayesian reaction optimization as a tool for chemical synthesis. 590(7844):89–96. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-021-03213-y. URL https://www.nature.com/articles/s41586-021-03213-y.

Shved, A. S., Ocampo, B. E., Burlova, E. S., Olen, C. L., Rinehart, N. I., and Denmark, S. E. molli: A general purpose python toolkit for combinatorial small molecule library generation, manipulation, and feature extraction. *Journal of Chemical Information and Modeling*, 64(21):8083–8090, October 2024. ISSN 1549-960X. doi: 10.1021/acs.jcim.4c00424. URL http://dx.doi.org/10.1021/acs.jcim.4c00424.

Sin, J. W., Chau, S. L., Burwood, R. P., Püntener, K., Bigler, R., and Schwaller, P. Highly parallel optimisation of chemical reactions through automation and machine intelligence. 16(1):6464. ISSN 2041-1723. doi: 10.1038/s41467-025-61803-0. URL https://www.nature.com/articles/s41467-025-61803-0.

Souza, L. W., Ricke, N. D., Chaffin, B. C., Fortunato, M. E., Jiang, S., Soylu, C., Caya, T. C., Lau, S. H., Wieser, K. A., Doyle, A. G., and Tan, K. L. Applying active learning toward building a generalizable model for ni-photoredox cross-electrophile coupling of aryl and alkyl bromides. *Journal of the American Chemical Society*, 147(22):18747–18759, May 2025. ISSN 1520-5126. doi: 10.1021/jacs.5c02218. URL http://dx.doi.org/10.1021/jacs.5c02218.

Taylor, C. J., Pomberger, A., Felton, K. C., Grainger, R., Barecka, M., Chamberlain, T. W., Bourne, R. A., Johnson, C. N., and Lapkin, A. A. A Brief Introduction to Chemical Reaction Optimization. 123(6):3089–3126. ISSN 0009-2665, 1520-6890. doi: 10.1021/acs.chemrev.2c00798. URL https://pubs.acs.org/doi/10.1021/acs.chemrev.2c00798.

Taylor, C. J., Felton, K. C., Wigh, D., Jeraal, M. I., Grainger, R., Chessari, G., Johnson, C. N., and Lapkin, A. A. Accelerated chemical reaction optimization using multi-task learning. *ACS Central Science*, 9(5):957–968, April 2023. ISSN 2374-7951. doi: 10.1021/acscentsci.3c00050. URL http://dx.doi.org/10.1021/acscentsci.3c00050.

Torres, J. A. G., Lau, S. H., Anchuri, P., Stevens, J. M., Tabora, J. E., Li, J., Borovika, A., Adams, R. P., and Doyle, A. G. A Multi-Objective Active Learning Platform and Web App for Reaction Optimization. 144(43):19999–20007. ISSN 0002-7863, 1520-5126. doi: 10.1021/jacs.2c08592. URL https://pubs.acs.org/doi/10.1021/jacs.2c08592.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Huggingface's transformers: State-of-the-art natural language processing, 2019. URL https://arxiv.org/abs/1910.03771.

Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., Dong, G., Wei, H., Lin, H., Tang, J., Wang, J., Yang, J., Tu, J., Zhang, J., Ma, J., Xu, J., Zhou, J., Bai, J., He, J., Lin, J., Dang, K., Lu, K., Chen, K., Yang, K., Li, M., Xue, M., Ni, N., Zhang, P., Wang, P., Peng, R., Men, R., Gao, R., Lin, R., Wang, S., Bai, S., Tan, S., Zhu, T., Li, T., Liu, T., Ge, W., Deng, X., Zhou, X., Ren, X., Zhang, X., Wei, X., Ren, X., Fan, Y., Yao, Y., Zhang, Y., Wan, Y., Chu, Y., Liu, Y., Cui, Z., Zhang, Z., and Fan, Z. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

Zhang, J., Sugisawa, N., Felton, K. C., Fuse, S., and Lapkin, A. A. Multi-objective bayesian optimisation using q -noisy expected hypervolume improvement (q nehvi) for the schotten–baumann reaction. *Reaction Chemistry amp; Engineering*, 9(3):706–712, 2024. ISSN 2058-9883. doi: 10.1039/d3re00502j. URL http://dx.doi.org/10.1039/D3RE00502J.

Zhang, T. Y. Process Chemistry: The Science, Business, Logic, and Logistics. 106(7):2583–2595. ISSN 0009-2665. doi: 10.1021/cr040677v. URL https://doi.org/10.1021/cr040677v.

Zhao, H., Ravn, A. K., Haibach, M. C., Engle, K. M., and Johansson Seechurn, C. C. C. Diversification of pharmaceutical manufacturing processes: Taking the plunge into the non-pgm catalyst pool. *ACS Catalysis*, 14(13):9708–9733, 2024. ISSN 2155-5435. doi: 10.1021/acscatal.4c01809. URL http://dx.doi.org/10.1021/acscatal.4c01809.

Zimmermann, Y., Bazgir, A., Al-Feghali, A., Ansari, M., Bocarsly, J., Brinson, L. C., Chiang, Y., Circi, D., Chiu, M.-H., Daelman, N., Evans, M. L., Gangan, A. S., George, J., Harb, H., Khalighinejad, G., Takrim Khan, S., Klawohn, S., Lederbauer, M., Mahjoubi, S., Mohr, B., Mohamad Moosavi, S., Naik, A., Beste Ozhan, A., Plessers, D., Roy, A., Schöppach, F., Schwaller, P., Terboven, C., Ueltzen, K., Wu, Y., Zhu, S., Janssen, J.,

Li, C., Foster, I., and Blaiszik, B. 32 examples of llm applications in materials science and chemistry: towards automation, assistants, agents, and accelerated scientific discovery. *Machine Learning: Science and Technology*, 6(3):030701, 2025. ISSN 2632-2153. doi: 10.1088/2632-2153/ae011a. URL http://dx.doi.org/10.1088/2632-2153/ae011a.

A. Supplemental main text benchmarking studies

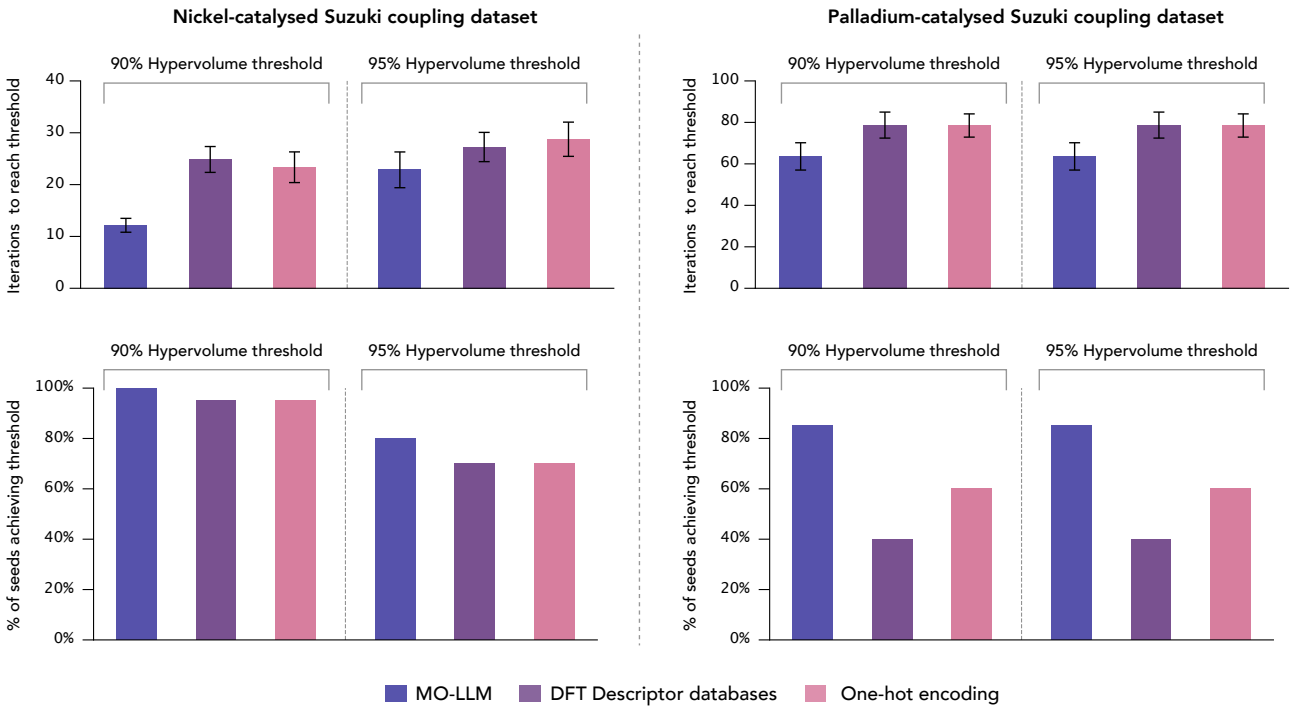

Figure 3. Sequential low-data optimisation benchmarks comparing different representation methods (MO-LLM, DFT descriptor databases, and one-hot encoding) on the nickel-catalysed and palladium-catalysed Suzuki coupling datasets. Top panels: mean number of iterations required to reach 90% and 95% hypervolume thresholds (lower is better, error bars indicate standard error across 20 random seeds). Bottom panels: percentage of optimisation runs (initialised with different random seeds) reaching each threshold within the allocated experimental budget (higher is better).

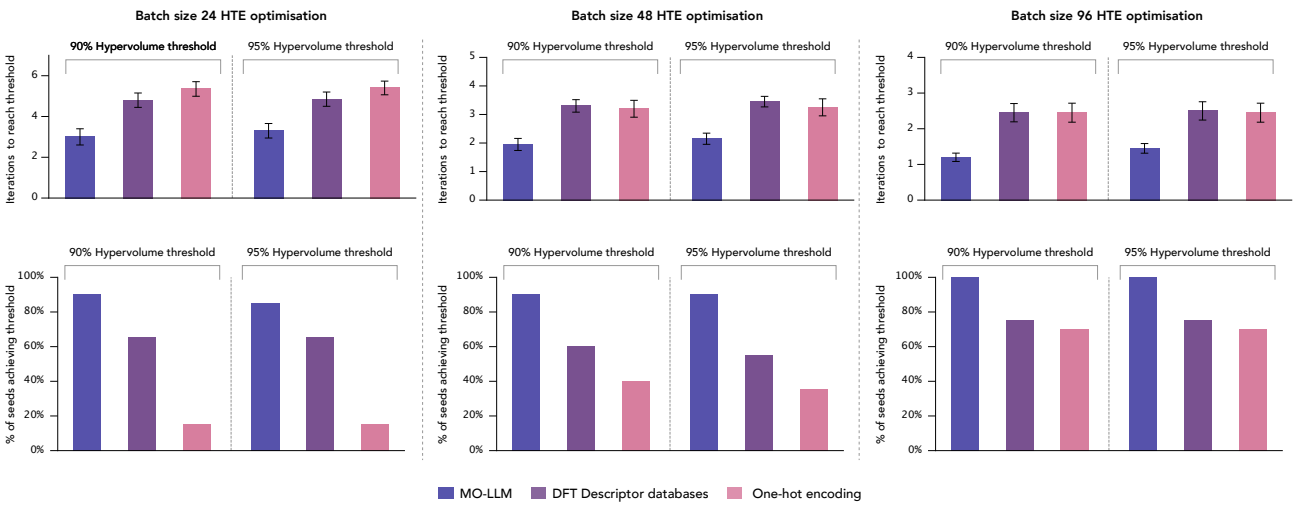

Figure 4. High-throughput experimentation (HTE) optimisation benchmarks on the palladium-catalysed sulfonamide coupling dataset, comparing MO-LLM, DFT descriptor databases, and one-hot encoding across 24-well (left), 48-well (middle), and 96-well (right) plate batch sizes. Top panels: mean number of plate iterations required to reach 90% and 95% hypervolume thresholds (lower is better, error bars indicate standard error across 20 random seeds). Bottom panels: percentage of optimisation runs (initialised with different random seeds) reaching each threshold within the allocated experimental budget (higher is better).

B. Influence of different pre-trained models and pooling strategies

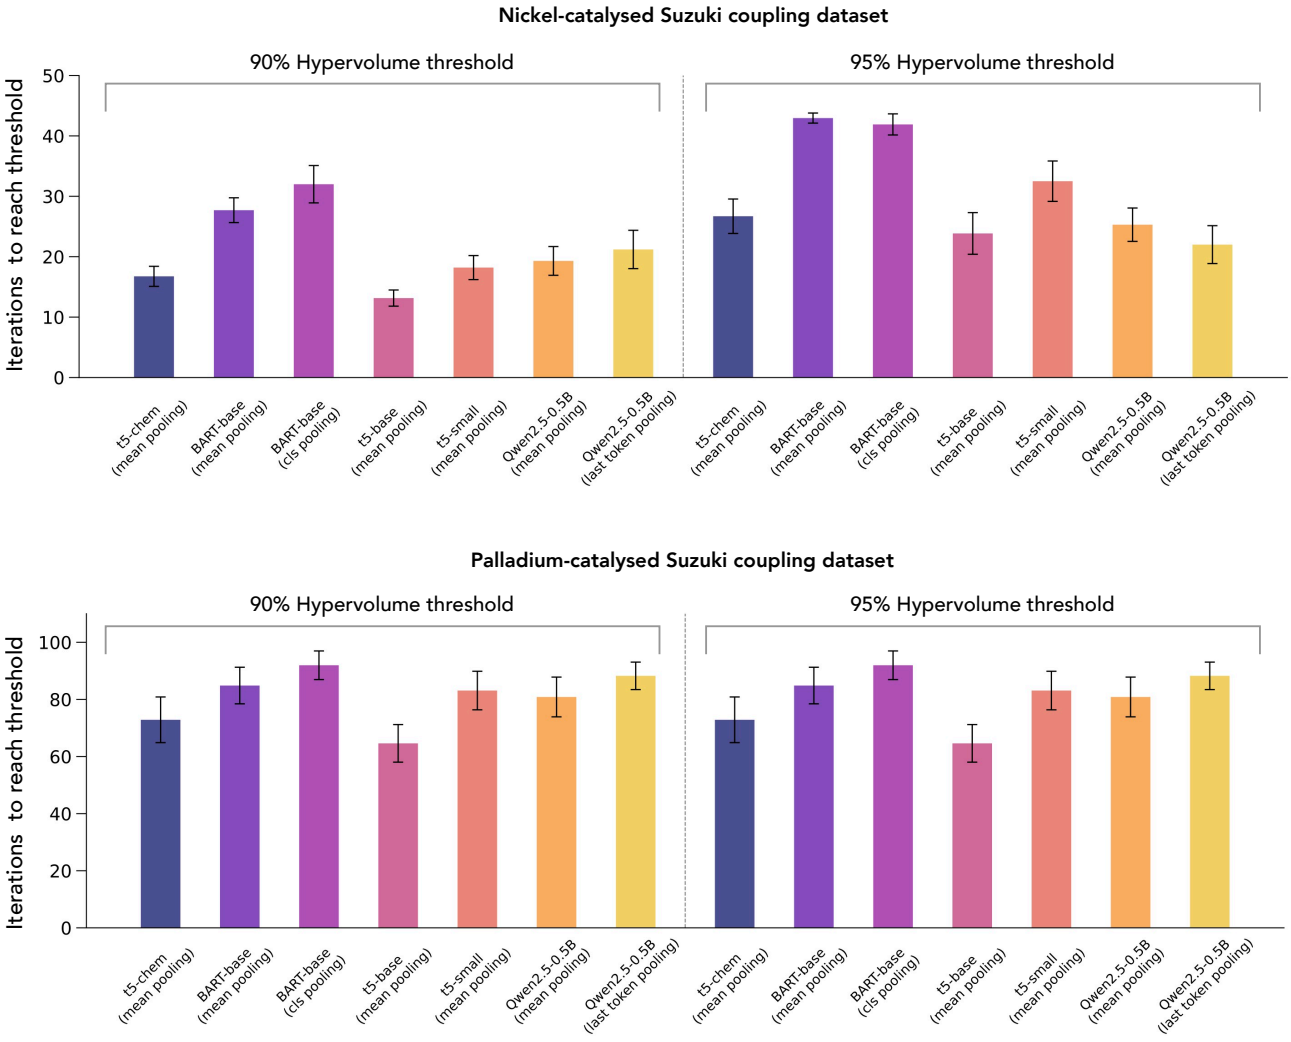

Figure 5. Ablation studies over pre-trained language model architectures and pooling strategies on the nickel-catalysed (top) and palladium-catalysed (bottom) Suzuki coupling datasets. Mean number of iterations required to reach 90% and 95% hypervolume thresholds are compared across T5-chem (mean pooling), BART-base (mean and CLS pooling), T5-base and T5-small (mean pooling), and Qwen2.5-0.5B (mean and last token pooling). Error bars indicate standard error across 20 random seeds. T5-base with mean pooling consistently achieved practical hypervolume thresholds within the fewest iterations across all datasets and thresholds.

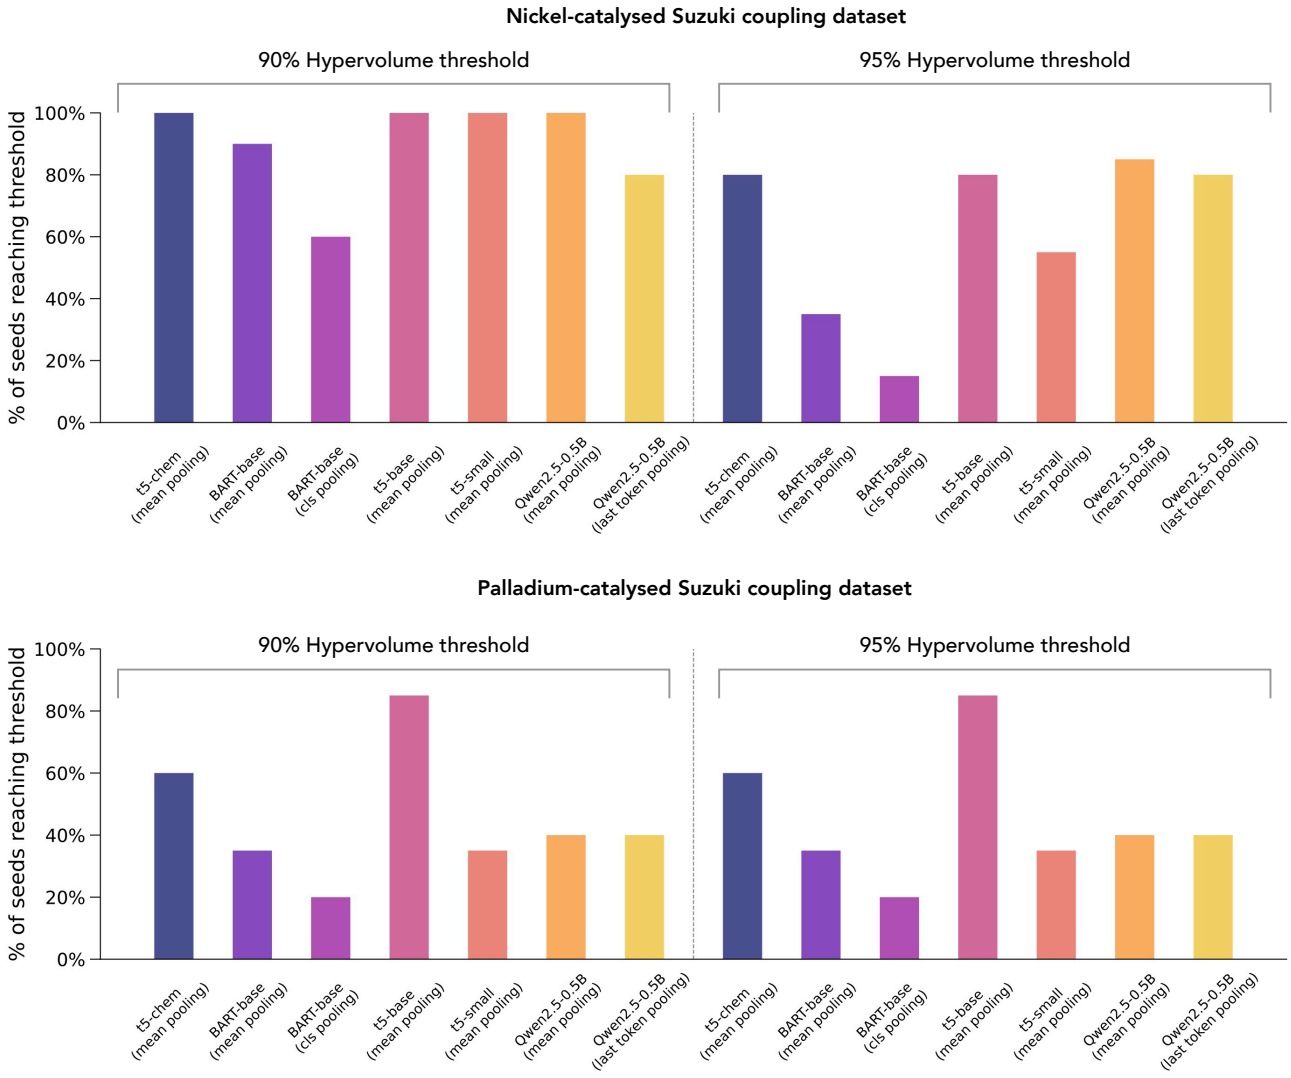

Figure 6. Percentage of optimisation runs (initialised with 20 different random seeds) reaching 90% and 95% hypervolume thresholds within the allocated experimental budget, across pre-trained language model architectures and pooling strategies on the nickel-catalysed (top) and palladium-catalysed (bottom) Suzuki coupling datasets. Models and pooling strategies are as described in Section 2.2.

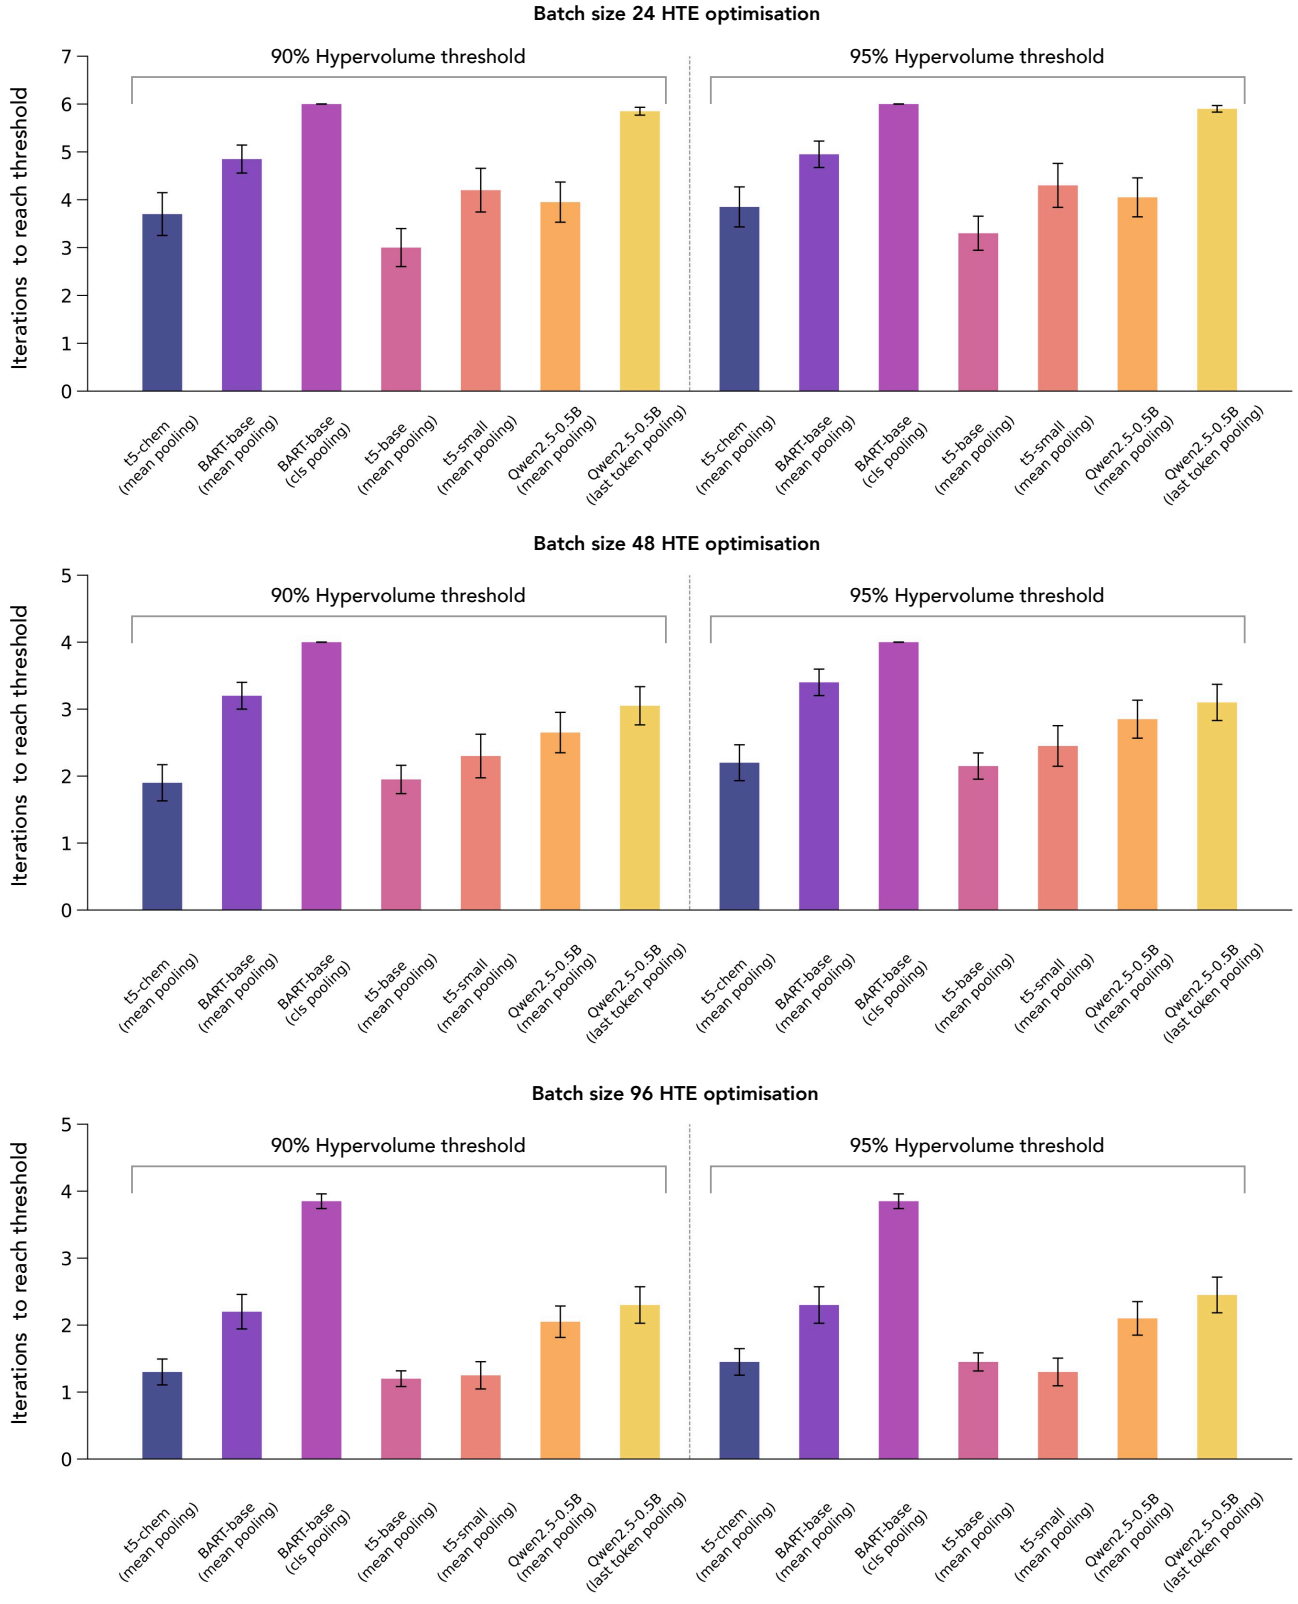

Figure 7. Ablation studies over pre-trained language model architectures and pooling strategies for high-throughput experimentation (HTE) optimisation on the palladium-catalysed sulfonamide coupling dataset, across 24-well (top), 48-well (middle), and 96-well (bottom) plate batch sizes. The bar charts show the mean number of plate iterations required to reach 90% and 95% hypervolume thresholds. Models and pooling strategies are as described in Section 2.2. Error bars indicate standard error across 20 random seeds.

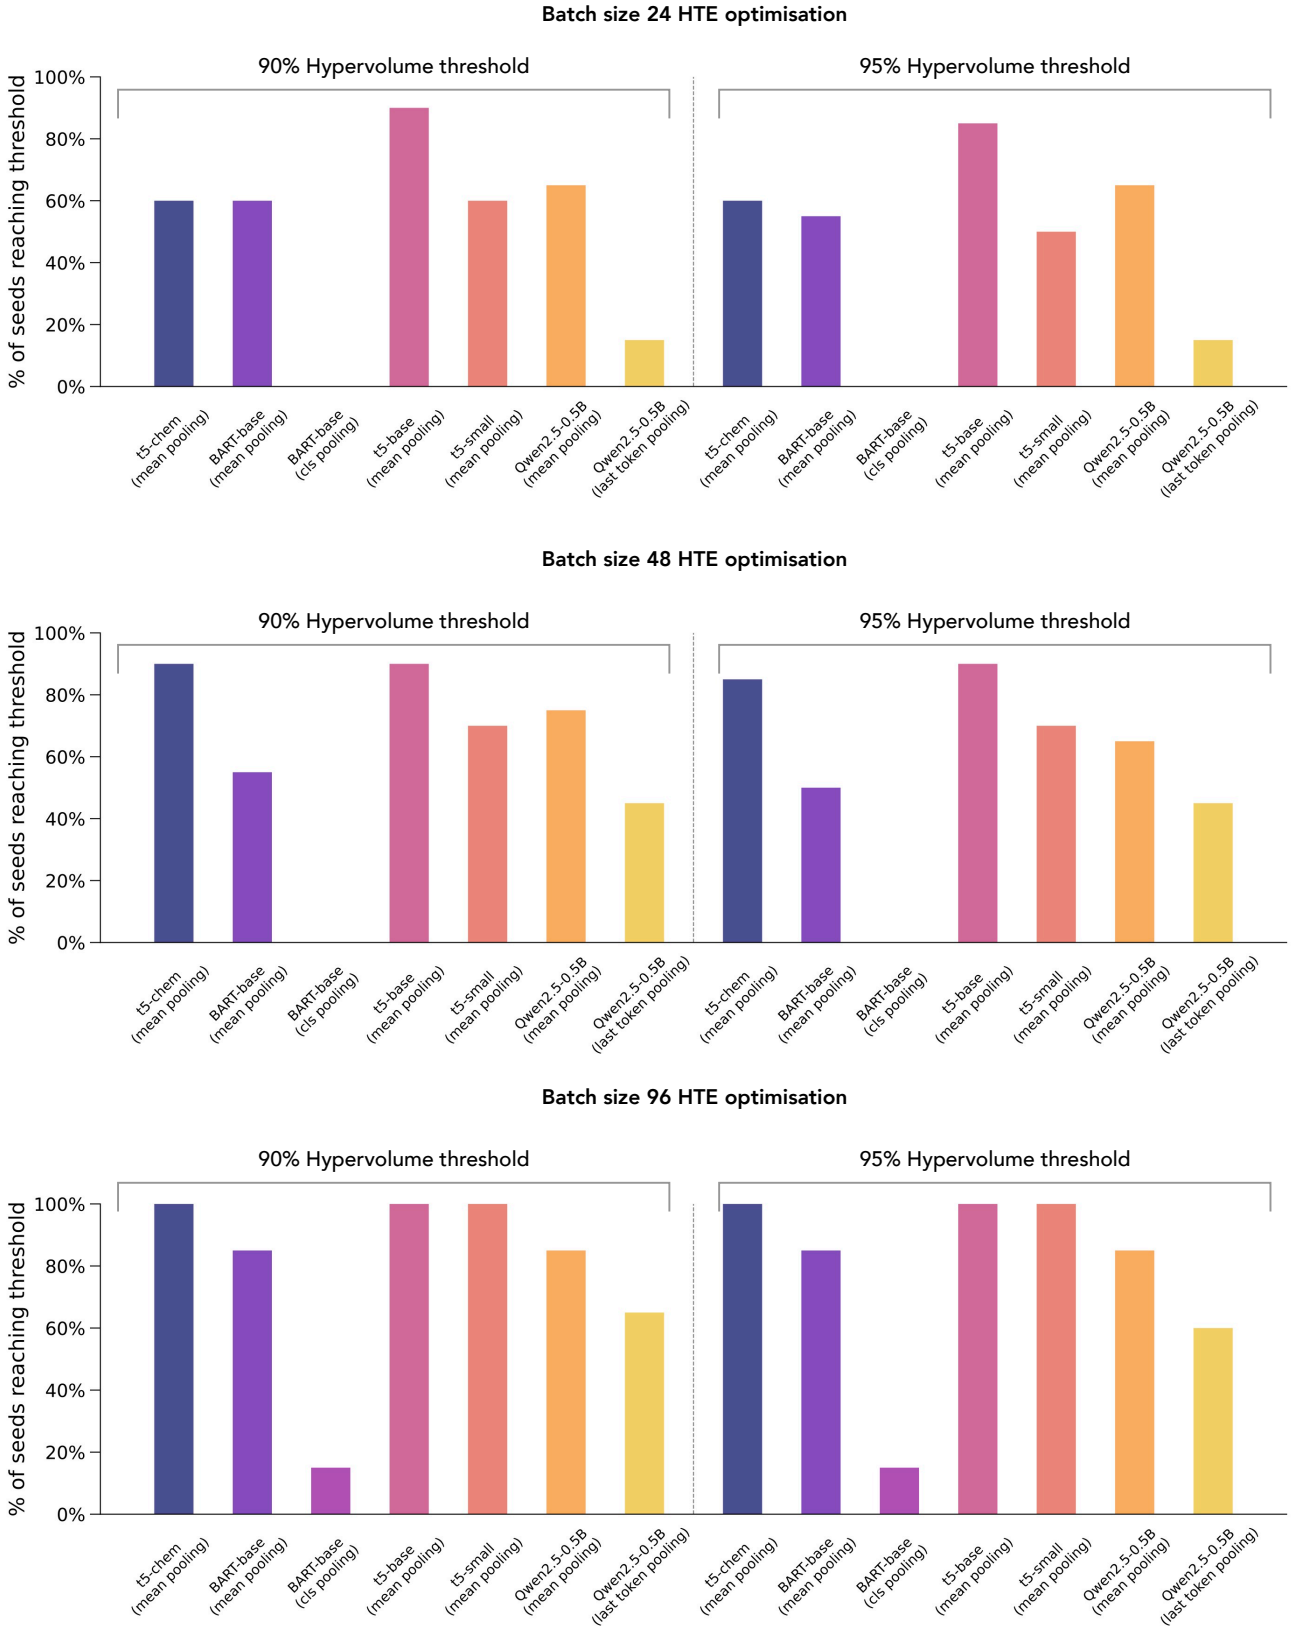

Figure 8. Percentage of optimisation runs (initialised with 20 different random seeds) reaching 90% and 95% hypervolume thresholds within the allocated experimental budget, across pre-trained language model architectures and pooling strategies for high-throughput experimentation (HTE) optimisation benchmarks on the palladium-catalysed sulfonamide coupling dataset, across 24-well (top), 48-well (middle), and 96-well (bottom) plate batch sizes. Models and pooling strategies are as described in Section 2.2.