Supporting Information:

Distilling and exploiting quantitative insights from Large Language Models for enhanced Bayesian optimization of chemical reactions

Roshan Patel[1], Mingxuan Li[2], Chin-Fei Chang[1], Louis De Lescure[1], Paul Chauvin[3], Alan Cherney[1], Saeed Moayedpour[2], Sven Jager[4], and Yasser Jangjou [*1]

[1]CMC Synthetics Platform, Sanofi, 350 Water St, Cambridge, MA, 02141, USA
[2]Digital R&D, Sanofi, 450 Water St, Cambridge, MA, 02141, USA
[3]Digital R&D, Sanofi, 58-60 avenue de la Grande Armée, Paris, France
[4]Digital R&D, Sanofi, Frankfurt 65929, Germany

S1 Selection and Parameterization of the Step Function $p(n)$

For our work, we constructed p(n) as a step function parameterized by three values:

$$p(n) = \begin{cases} v_1 & \text{if } n \leq c_1 \\ v_2 & \text{if } c_1 < n \leq 40 \\ 0 & \end{cases}$$

To optimize these three parameters, we used 300 iterations of the Tree-Structured Parzen Estimator [3] implemented by the Hyperopt package [2] over the parameter space: $v_1 \in [70, 95]$, $v_2 \in [0, 70]$, $c_1 \in [1, 39]$ and select the parameter set that gives the best performance among datasets BH1-5 relative to the baseline EI. Specifically, the minimization objective was computed using the following formula:

$$a(v_1, v_2, c_1) = \sum_{i=1}^{5} (\bar{n}_{\text{max,EI-LLM,BH}i}(v_1, v_2, c_1) - \bar{n}_{\text{max,EI,BH}i}) - 50 \cdot \mathbb{I}\left[\forall i,\ \bar{n}_{\text{max,EI-LLM,BH}i} < \bar{n}_{\text{max,EI,BH}i}\right] \tag{1}$$

where $\bar{n}_{\text{max,EI-LLM,BH}i}(v_1, v_2, c_1)$ is the number of experiments needed to observe the maximum yield (averaged over 50 trials) using the EI w/LLM acquisition function for dataset BHi, $\bar{n}_{\text{max,EI,BH}i}$ is the corresponding quantity obtained using the standard EI acquisition function, and $\mathbb{I}[\cdot]$ denotes the indicator function. The indicator term applies a fixed bonus of -50 whenever the EI w/LLM acquisition function requires fewer experiments than standard EI across all five datasets used for parameter optimization(BH1–BH5). This bonus term was introduced to explicitly favor parameter settings that consistently improved optimization performance across all datasets to emphasize robustness in the designed $p(n)$.

Optimization yields parameters $v_1 = 85\%$, $v_2 = 15\%$, and $c_1 = 30$ giving the $p(n)$ shown in Figure S1. We use this $p(n)$ for all BO trials employing the EI w/LLM acquisition function in Figure 3 and Figure 4. The good performance of LLM-EI on the DA dataset shown in Figure 3 and the three additional AC 1-3 dataset in Figure 4 (not used for parameter optimization) serves as validation of the $p(n)$ optimized in our study.

*Corresponding author: Yasser.Jangjou@sanofi.com

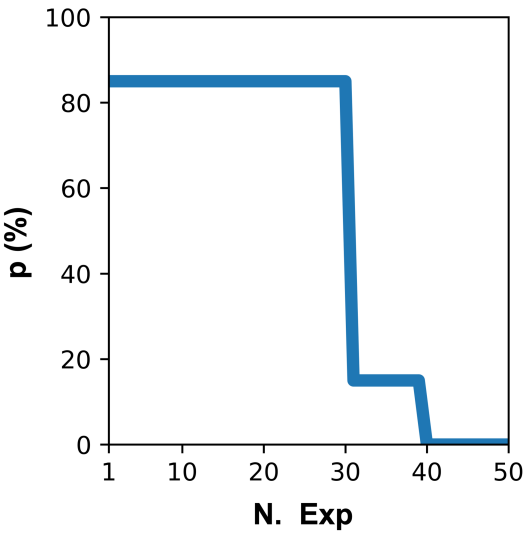

Figure S1: Optimized $p(n)$ used for the LLM-EI acquisition function. The plot shows the cutoff percentile p% (y-axis) for outputs of g(x) as a function of iterations of Bayesian optimization (x-axis). For example, according to this function and the definition of $\pi(x)$ presented in section 2.2.3, the first 30 experiments are selected from the subset experiments with preference values above the 85 percentile of g(x) for $x \in X$.

S2 Supplementary Figures

To understand the performance of LLMs in reasoning the relation of chemical reaction conditions and yield, four LLMs (Sonnet-3.5, Sonnet-3, haiku-3 and GPT-4) were applied in the yield preference survey. Each datasets has 1,000 question pairs and the accuracy for each LLM is shown in Figure S2. Accuracy of a full-length survey for each dataset using Sonnet-3.5 is also presented (purple dot line).

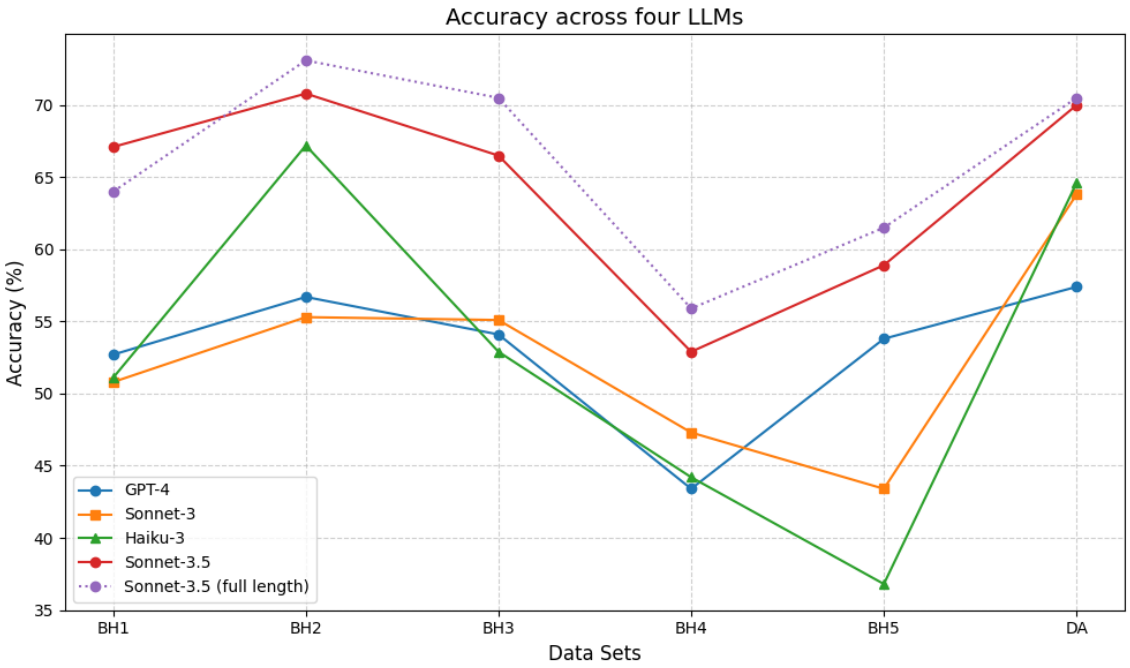

Figure S2: Comparison of LLMs performance in chemical reaction yield reasoning.

Figure S3 shows an example of typical reasoning used by the LLM to answer survey questions. This example is a randomly selected question from the survey generated for the DA dataset.

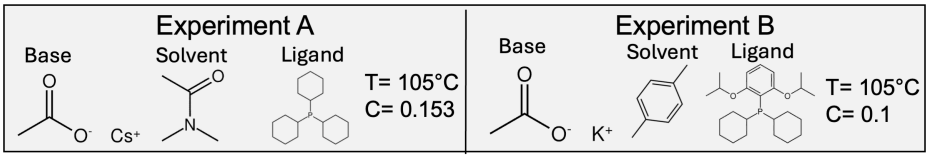

Example Prompt
"Task: For the following reaction, predict which experiment setup leads to a higher yield and output the response (A or B) along with reasoning in JSON format. Task reaction: "Direct Arylation" For any reaction, an additive, a base, and a palladium-ligand are present.

Setup A:
Base SMILES: O=C([O-])C.[Cs+]
Solvent SMILES: CC(N(C)C)=O
Ligand SMILES: P(C1CCCCC1)(C2CCCCC2)C3CCCCC3
Temperature: 105°C
Concentration: 0.153 M

Setup B:
Base SMILES: O=C([O-])C.[K+]
Solvent SMILES: CC1=CC=C(C)C=C1
Ligand SMILES: CC(OC1=C(P(C2CCCCC2)C3CCCCC3)C(OC(C)C)=CC=C1)C
Temperature: 105°C
Concentration: 0.1 M

Output: Reaction setup (A or B) with higher yield and the reasoning in a JSON object with 'Setup' and 'reasoning' keys."

Example Response
"Setup A is likely to lead to a higher yield for the given reaction due to the following reasons:
1- Base: Cesium carbonate (Cs2CO3) in Setup A is generally a stronger base than potassium carbonate (K2CO3) in Setup B. Stronger bases are often more effective in promoting cross-coupling reactions.
2- Ligand: The ligand in Setup A (triphenylphosphine) is a common and effective ligand for many cross-coupling reactions. It's simpler structure may provide better stability and catalytic activity compared to the more complex ligand in Setup B.
3- Solvent: Dimethylacetamide (DMAc) in Setup A is a polar aprotic solvent that is often used in cross-coupling reactions. It can effectively dissolve both organic and inorganic compounds, which is beneficial for this type of reaction. Toluene in Setup B is less polar and may not be as effective.
4- Concentration: Setup A has a higher concentration (0.153 M) compared to Setup B (0.1 M), which could lead to faster reaction rates and potentially higher yields.
5- Temperature: Both setups have the same temperature, so this factor does not differentiate between them.
6- Overall, the combination of a stronger base, effective ligand, suitable solvent, and higher concentration in Setup A is more likely to result in a higher yield for this cross-coupling reaction."

Figure S3: Example of LLM reasoning in answering survey

Figure S4 provides a comparison between two methods aiming to extract quantitative insights from the LLM. Panel a shows the result of zero-shot LLM regression for the DA dataset, i.e., the LLM was now described the conditions of a single experiment (in the same way that was done for answering the survey) but instead prompted to provide the predicted yield instead of a preference. Panel b shows the utility function $g(x)$ outputs. The utility function outputs show a modest correlation to the true experimentally measured yield, whereas the LLM-predicted yields do not.

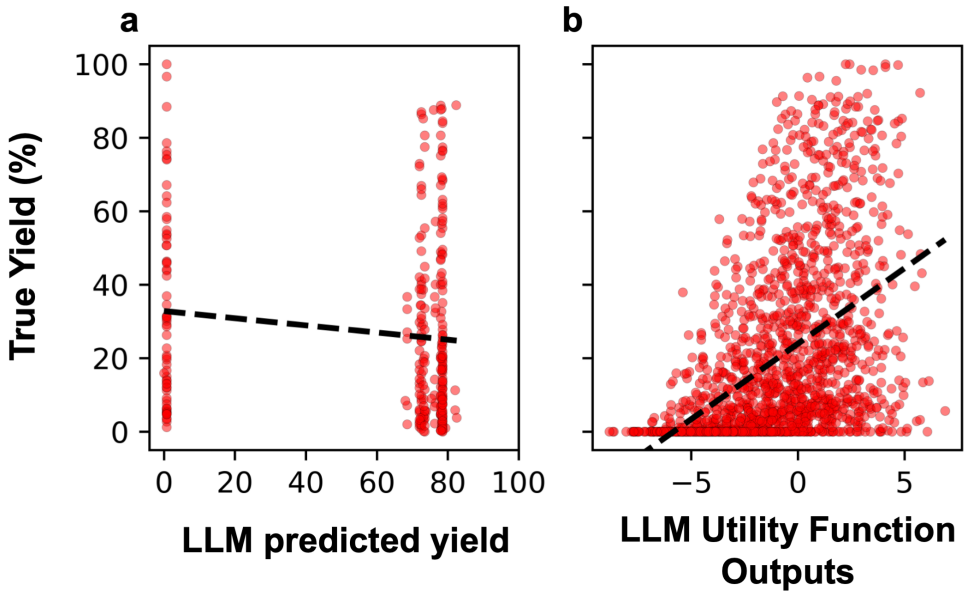

Figure S4: Comparison between simple zero-shot regression (panel a) and preference learning on survey (answered by zero-shot LLM) for the DA dataset (panel b)

Figure S5 shows the mean (and standard error from n=50 trials) number of experiments needed to run to observe 99% of the maximum yield for a given dataset and acquisition function. The x-axis shows dataset names ranging from Buchwald-Hartwig (BH) 1 to 5 and Direct Arylation (DA) (left to right), and the y-axis shows the number of experiment to reach 99% maximum.

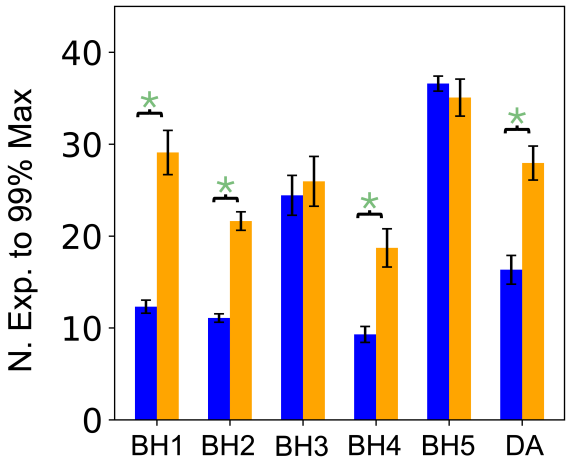

Figure S5: Comparing the average number of experiments needed to identify experimental conditions that give 99% of the maximum observed yield using the LLM-EI (blue) and EI (yellow) acquisition functions for datasets BH1-BH5 and DA. Refer to the caption of Figure 3 for additional details regarding significance testing.

S3 Gaussian Process Surrogate Model Developed by Shields et al.[6]

The Gaussian process surrogate model used in this work follows the formulation introduced by Shields et al.[6]. The reaction outcomes were standardized to zero mean and unit variance and modeled by a constant mean Gaussian process with a Matérn kernel of shape parameter $v = 5/2$. The kernel function is given by:

$$k_{Matérn52}(r) = \alpha(1 + \frac{\sqrt{5}r}{l} + \frac{5r^2}{3l^2}) \exp(-\frac{\sqrt{5}r}{l}) \tag{2}$$

where α is an output scale parameter, $r = \sqrt{\sum_{i=1}^{n}(\mathbf{x}_{1i} - \mathbf{x}_{2i})^2}$ is the distance between two input points ($\mathbf{x}_1$ and $\mathbf{x}_2$), and l is the characteristic length scale parameter.

In addition, automatic relevance determination [7] was employed by learning a separate length scale for each input dimension. Gamma priors favoring longer length scales were placed on the length-scale parameters to regularize the model in the low-data regime. Output variance and observation noise variance were learned during training. All kernel and likelihood hyperparameters were optimized by maximizing the log marginal likelihood using stochastic gradient-based optimization, as implemented in GPyTorch[5]

S4 Pairwise Gaussian Process Preference Model

The Preference learning was performed using a Pairwise Gaussian process (PairwiseGP) model provided in the BoTorch[1] library, which is an approach introduced by Chu and Ghahramani[4]. A latent function $g(x)$ is assumed, and observed preferences are modeled as comparisons between pairs of inputs, which are the reaction conditions preference generated by the LLMs survey. Specifically, when one reaction condition x_j is preferred than another x_k, the model assumes that the latent value of the preferred condition is larger $g(x_j) > g(x_k)$.

In this present work, the kernel is set as the default Scaled RBF kernel defined in PairwiseGP, and the model employs the `PairwiseLikelihood`, $\frac{g(x_j)-g(x_k)}{\sqrt{2}\sigma}$, where the σ is implicitly set to 1 as the

function is scaled by the implemented kernel. Posterior inference over the latent function was carried out using a Laplace approximation, as implemented via `PairwiseLaplaceMarginalLogLikelihood` in BoTorch. This approximation enables training of the preference model by optimizing a marginal log likelihood objective that accounts for uncertainty in pairwise comparison data.

References

[1] M. Balandat et al. "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization". In: *Advances in Neural Information Processing Systems 33*. 2020.

[2] J. Bergstra, D. Yamins, and D. Cox. "Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures". In: *Proceedings of the 30th International Conference on Machine Learning*. Ed. by S. Dasgupta and D. McAllester. Vol. 28. Proceedings of Machine Learning Research 1. Atlanta, Georgia, USA: PMLR, 17–19 Jun 2013, pp. 115–123.

[3] J. Bergstra et al. "Algorithms for Hyper-Parameter Optimization". In: *Advances in Neural Information Processing Systems*. Ed. by J. Shawe-Taylor et al. Vol. 24. Curran Associates, Inc., 2011.

[4] W. Chu and Z. Ghahramani. "Preference learning with Gaussian processes". In: *Proceedings of the 22nd international conference on Machine learning* (2005).

[5] J. Gardner et al. "Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration". In: *Advances in neural information processing systems* 31 (2018).

[6] B. J. Shields et al. "Bayesian reaction optimization as a tool for chemical synthesis". en. In: *Nature* 590.7844 (Feb. 2021), pp. 89–96.

[7] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. Vol. 2. 3. MIT press Cambridge, MA, 2006.