White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/182950/

Version: Accepted Version

Article:
Mueller, P, Clayton, AD, Manson, J et al. (2022) Automated Multi-Objective Reaction Optimisation: Which Algorithm Should I Use? Reaction Chemistry and Engineering. ISSN: 2058-9883

https://doi.org/10.1039/D1RE00549A

© The Royal Society of Chemistry 2022. This is an author produced version of an article published in Reaction Chemistry & Engineering. Uploaded in accordance with the publisher's self-archiving policy.

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Please do not adjust margins

Reaction Chemistry & Engineering

ARTICLE

Received 00th January 20xx,
Accepted 00th January 20xx

DOI: 10.1039/x0xx00000x

www.rsc.org/

Automated Multi-Objective Reaction Optimisation: Which Algorithm Should I Use?

Pia Mueller,[†a] Adam D. Clayton,[†a] Jamie Manson,[a] Samuel Riley,[a] Oliver S. May,[bc] Norman Govan,[b] Stuart Notman,[b] Steven V. Ley,[c] Thomas W. Chamberlain[a] and Richard A. Bourne*[a]

Multi-objective optimisation algorithms (MOOAs) are, of increasing interest for the efficient optimisation of chemical processes. However, an algorithms performance can vary on a case-by-case basis, depending on the complexity of the search space and the nature of the underlying response surfaces. This makes appropriate algorithm selection for chemical reaction optimisation a challenging problem. An open-source reaction simulator has been developed, which enables the performance of multi-objective algorithms to be benchmarked against a series of chemistry-inspired test problems. The performance of four different MOOAs were compared, including three state-of-the-art Bayesian optimisation algorithms, and their ability to optimise different types of systems quantified using the hypervolume metric. In general, EIMEGO was found to achieve the highest hypervolume in the lowest number of experiments and was only outperformed by TSEMO in cases with three objectives. To verify the simulated results, EIMEGO and TSEMO were tested experimentally using a three-objective optimisation of methyl phenyl sulfide oxidation. Both algorithms successfully identified a trade-off between conversion, selectivity, and productivity with respect to the desired sulfoxide. In this case, TSEMO outperformed EIMEGO in terms of hypervolume, which was in agreement with the simulated results.

Introduction

With the current digital transformation in manufacturing, known as Industry 4.0, there has been a rise in the digitalisation of chemistry. Intelligent systems fuelled by data and machine learning enable the automation of complex tasks, such as route design and reaction discovery, as well as more routine experimentation, such as synthesis, reaction screening and optimisation studies.[1–5] Specifically, 'self-optimising' systems utilise machine learning algorithms to define the next set of reaction conditions to explore, based on the results of the previous experiments.[6] Reaction optimisation is considered to be an expensive-to-evaluate problem, with staff time and high-value materials contributing to a significant expense. Therefore, it is imperative that self-optimising systems are able to find the global optimum in as few experiments as possible, whilst maximising the amount of information gained per experiment. This has recently led to the application of Bayesian optimisation algorithms, which aim to balance the trade-off between exploration and exploitation.[7] This approach has also been used to enable multi-objective reaction optimisation, which has been shown to save time and resources by simultaneously optimising multiple performance criteria.[8] Notably, these identify the trade-off curve (Pareto front) between conflicting objectives (e.g. economic vs. environmental), which provides valuable information during the development of chemical processes.[9,10]

The end-user of these systems is typically a chemist, who in general does not have an extensive knowledge of programming and algorithm design. However, the efficiency of the optimisations is dependent on the algorithm selected on a case-by-case basis. So, how does a chemist know which algorithm to use and when? This question does not have a single answer, as the performance of algorithms can vary depending on the type of problem.[11] Furthermore, it would be unrealistic to compare multiple algorithms using an experimental platform, as this would be very expensive and time-consuming. Therefore, in computer science, algorithm performance is compared on a series of different *in silico* test problems which are designed to have different characteristics. However, these test problems are mathematical functions, which do not relate to chemical

[a]Institute of Process Research and Development, School of Chemistry & School of Chemical and Process Engineering, University of Leeds, LS2 9JT, UK. E-mail: r.a.bourne@leeds.ac.uk

[b]Defense Science and Technology Lab, Porton Down, Salisbury, SP4 0JQ, UK

[c]Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge, CB2 1EW, UK

† Joint first authors who contributed equally to this work and the preparation of themanuscript.

This journal is © The Royal Society of Chemistry 20xx

Please do not adjust margins

reaction applications. A series of chemistry inspired test problems based on reaction models would be more informative. This approach was recently reported by Lapkin *et al.*, where an open-source framework (Summit) was produced to compare algorithms on chemically-motivated benchmarks.[12] Herein, we report six test problems designed to compare multi-objective optimisation algorithms. The suitability of each algorithm for different types of chemical reactions is identified.

Additionally, a comparison of two algorithms is experimentally performed to validate the findings proposed by the theoretical models. The reaction of selective organic sulfide oxidation by hydrogen peroxide is used to showcase experimentally the differences in algorithms and the additional challenges. This model reaction is well studied due to its pharmaceutical value[13,14] but is furthermore of interest for simulation of chemical warfare agent degradation.[15,16] The reaction is accessible, safe and simplistic and therefore ideal to study experimentally. Whilst standard experimental studies have previously been able to find optima for one objective by changing one variable at a time,[17] self-optimized studies are able to correlate variables in a coherent functionality to multi-objectives.[18]

Multi-Objective Reaction Simulator

To compare the performance of multi-objective optimisation algorithms for optimising chemical systems, a kinetic-based reaction simulator was designed. Four reactions with known kinetic parameters (pre-exponential factors and activation energies) were chosen from the literature: (i) Van de Vusse reaction (VdV1); (ii) nucleophilic aromatic substitution (S_NAr1 & S_NAr2) between 2,4-difluoronitrobenzene and morpholine; (iii) isomerisation of lactose to lactulose (Lactose1); (iv) Paal-Knorr reaction (PK1 & PK2) between 2,5-hexanedione and ethanolamine (see SI for schemes).[19–22] These examples provided a good representation of non-competitive (iv) and competitive reactions, including competing parallel (i, ii & iii) and consecutive pathways (ii & iii). Although reactions (iii) and (iv) contain reversible reactions, the k_{-1} rate constants are negligible and were therefore omitted. Six test problems were formulated using reaction variable limits as constraints and different process metrics as objectives. Details and visual representations of the variable and objective space for each test problem are provided in the supporting information. A summary of the test problems including descriptions of the Pareto fronts is provided in Table 1. Each problem was designed to contain between 2-4 variables and 2-3 objectives, as higher dimensional problems are not often encountered with current self-optimising chemical platforms, due to the self-optimisation of more complex multi-step processes and mixed variable problems remaining in their infancy. Furthermore, the variable limits and objectives for each problem were selected to ensure that a diverse range of Pareto fronts were generated in terms of morphology, uniformity and continuity. This was important to ensure that the robustness and different capabilities of the algorithms could be benchmarked against the standardised test problems.[23]

The simulation procedure for each test problem is outlined below. Firstly, the pre-exponential factors, A, and activation energies, E_a, were used to calculate the rate constants, k, for each step in the reaction using the Arrhenius equation. The differential rate equations for each step were then solved using an ordinary differential equation (ODE) solver. This provided the percentage of each species in the reaction mixture under different sets of conditions, which were subsequently used to calculate the objectives for the given test problem. Random noise, at a level inherent with precise, automated experimental systems, was also included by applying a maximum absolute error of 0.25% and maximum relative error of 0.5% to the outputs.[24]

To compare the performance of the algorithms, the hypervolume was calculated after each iteration, where the hypervolume is defined as the volume between the current Pareto front and a reference point (i. e. larger hypervolume = better Pareto front). The hypervolume was calculated using a Monte-Carlo approximation, which determined the proportion of 100,000 random points in the objective space which were dominated by the current Pareto front, P. Formally, given n samples drawn uniformly from R^k, where k is the number of objectives, within the hyperrectangle $[P, R]$, with R being the reference point for the calculation, the following can be defined:[25,26]

$$H_{MC} = \frac{1}{n}\sum_{i=1}^{n} I(P \leq Y_i) \qquad (1)$$

Y_i is the ith objective vector generated from the Monte-Carlo sampling process. $I(\cdot)$ returns 1 if the argument is true and 0 if false. An illustrative example of the estimation procedure is provided in **Error! Reference source not found.**.

The utopian and anti-utopian points are defined as the combination of the best and the combination of the worst solutions for each objective respectively.

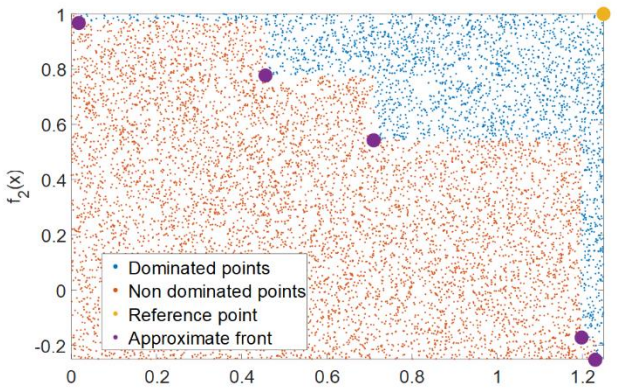

Figure 1. Illustrative example of Monte-Carlo based hypervolume approximation procedure. The hypervolume is the proportion of 100,000 random points in the objective space which are dominated by the current Pareto front.

These points were determined for the objective space of each problem by creating a superset of the non-dominated solutions from all runs across all algorithms. The reference

Please do not adjust margins

point for the objective space was then defined as the anti-utopian point shifted by 0.01 of the difference between the utopian and anti-utopian point.[27]

Table 1. Summary of the chemistry-inspired test problems designed using reported kinetic parameters.

Test problem	# Variables	# Objectives	Description of Pareto front
VdV1	2	2	Density of solutions fall away near to the Pareto front, which is non-uniformly distributed between linear and convex regions
S_NAr1	2	3	Optimal solutions follow a convoluted path through objective space with concave regions
S_NAr2	4	3	Convex, non-uniformly distributed Pareto front
Lactose1	2	2	Pareto front is a convex curve with many solutions
PK1	2	2	Pareto front is a convex curve with relatively few solutions
PK2	3	2	Pareto front consists of three discontinuous linear and concave regions

The code required to simulate each of these test problems is available online (https://github.com/adamc1994/MultiChem) and can be used to benchmark the performance of any multi-objective optimisation algorithm (MOOA). Herein, we use this simulator to benchmark the performance of five MOOAs for chemical reaction applications. As these applications require physical experiments to be conducted, which can take several tens of minutes and require expensive reagents, minimising the number of experiments required is paramount. Therefore, the optimisation performance is compared based on minimising the number of iterations required to achieve a high hypervolume.

Algorithm Performance

The performance of Thompson sampling efficient multi-objective optimisation (TSEMO), Pareto efficient global optimisation (ParEGO), NSGA-II and expected improvement matrix efficient global optimisation (EIM-EGO) were compared using the developed approach.[27–30] TSEMO, ParEGO and EIM-EGO are state-of-the-art Bayesian optimisation algorithms designed independently from each other for expensive-to-evaluate problems, thus making them well-suited for the optimisation of chemical reactions using experimental platforms. In contrast, NSGA-II is a computationally fast and elitist MOOA which generally requires many iterations, thus providing a suitable comparison for this study. Implementations of ParEGO, NSGA-II and EIM-EGO were all available in the platform for evolutionary multi-objective optimisation (PlatEMO) toolbox in MATLAB.[31] An implementation of TSEMO was available on GitHub, and was compared with experimental batches of one and four (batch sequential, BS-TSEMO) per iteration.[32] For the self-optimisation of chemical reactions, operating in batch sequential mode has the advantage of reducing the overall optimisation time in cases where the analysis time creates a bottleneck in the experiment.

The TSEMO, BS-TSEMO, ParEGO and EIM-EGO were initialised using a Latin hypercube (LHC, size = 20) for maximal experimental distribution.[33] Each algorithm had a function evaluation budget of 100 and was ran 20 times for each test problem to compare average performance. To account for the function evaluation budget, the NSGA-II population size and total number of generations were changed to 20 and 5 respectively.

Plots showing the average change in hypervolume throughout the optimisations are shown in **Error! Reference source not found.**. Furthermore, boxplots of the optimisation results after 60 function evaluations are displayed in **Error! Reference source not found.**, which provides a more detailed view at a practical number of experiments. From these results, the following observations and conclusions can be made:

- The NSGA-II algorithm has the lowest hypervolume after any number of function evaluations beyond the initial dataset in all examples apart from the Lactose1 test problem, where it has similar performance to ParEGO. Notably, NSGA-II requires less time per iteration (Table S1), but requires significantly more iterations to improve the hypervolume, and often fails to identify the complete trade-off curve in a practical number of experiments. Thus, surrogate model-based approaches are more suitable for expensive-to-evaluate problems.
- With the exception of NSGA-II, all algorithms exhibit a comparable performance for test problems S_NAr2 and PK1 after 60 experiments. Convergence of the algorithms on the Pareto front in less than 60 experiments suggests that these test problems are solved more easily. In general, ParEGO and EIM-EGO plateau at the maximum hypervolume in a lower number of experiments.
- The TSEMO and BS-TSEMO algorithms have a very similar performance based on the median and interquartile range of the hypervolume after 60 function evaluations. However, a notable difference occurs for the VdV1, SNAr2 and PK1 test problems, where the increase in hypervolume after the initial dataset is slower for BS-TSEMO. This can be attributed to the surrogate models been updated less frequently with new data when operating in batch sequential mode. It is observed from the results that this has a negative impact on algorithm performance in cases where the Pareto front is non-uniformly distributed. Therefore, although operating in batch sequential mode offers advantages in terms of reducing the overall self-optimisation time, this can lead to sub-optimal results.

Please do not adjust margins

- TSEMO was found to outperform ParEGO and NSGA-II in 5 out of 6 test problems. This is in agreement with previous benchmarking studies using mathematical test functions, where it outperformed ParEGO and NSGA-II in 6 out of 8 cases.[28]
- ParEGO produces a large range of hypervolumes in 3 of the 6 test problems (Lactose1, PK1, PK2), indicating that ParEGO is less robust compared to the other algorithms. ParEGO combines the objective functions linearly with random weightings. Therefore, this randomness could cause a spread in the results, which is emphasised in cases where the Pareto front is convex or discontinuous.
- EIM-EGO outperforms the other algorithms on 4 of the 6 test problems (VdV1, Lactose1, PK1, PK2) based on the median of the hypervolume after 60 function evaluations. Furthermore, the increase in hypervolume is significantly faster after the initial dataset compared to TSEMO for the VdV1 and PK2 test problems. This is likely caused by an increased emphasis on exploitation when using an expected improvement matrix compared to Thompson sampling.
- TSEMO outperforms EIM-EGO on test problems with 3 objectives (S_NAr1 and S_NAr2). This difference suggests that using an acquisition function which maximises hypervolume improvement is more efficient than one which maximises the Euclidean distance for problems with 3 objectives.

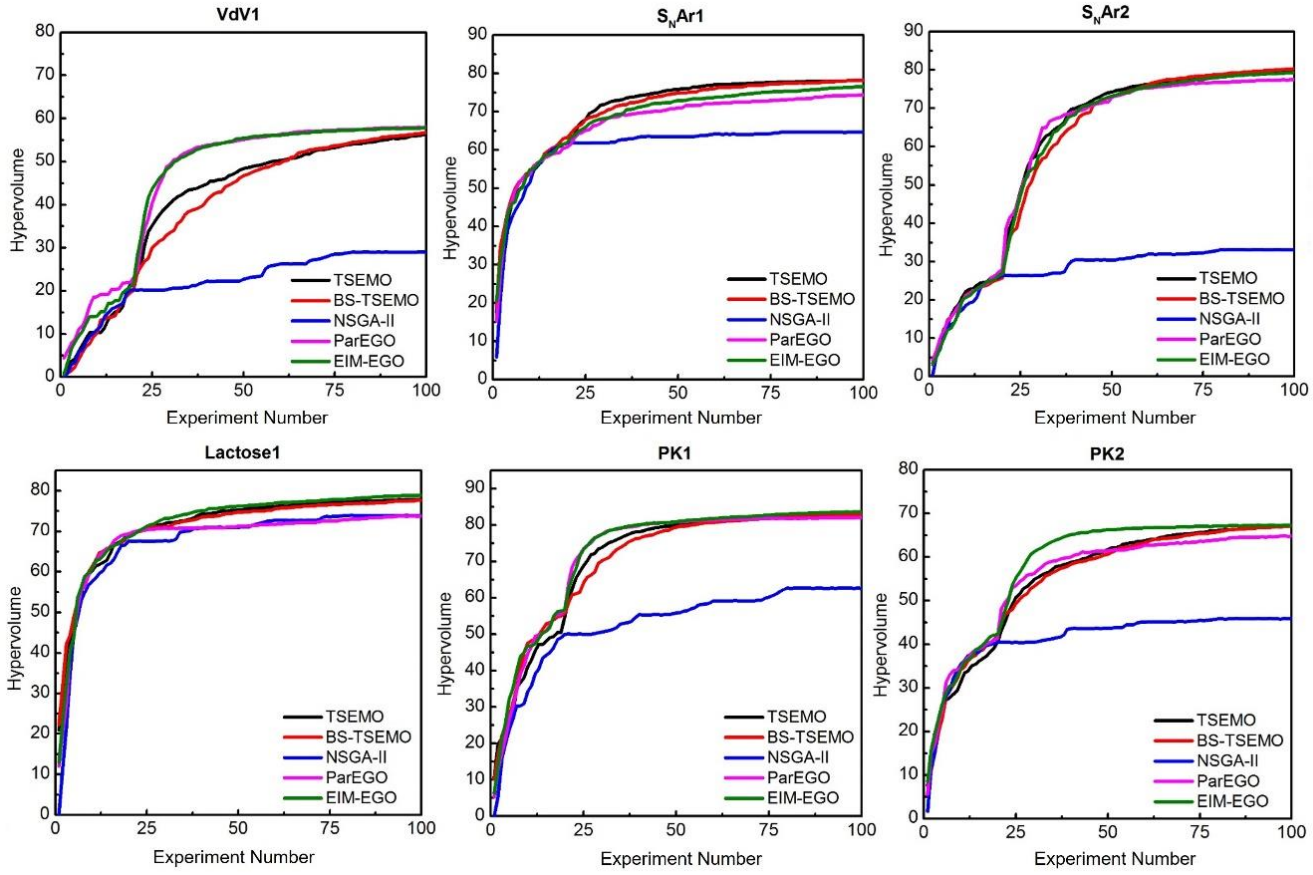

Figure 2. Plots showing the average change in hypervolume across 20 runs with 100 function evaluations each.

Please do not adjust margins

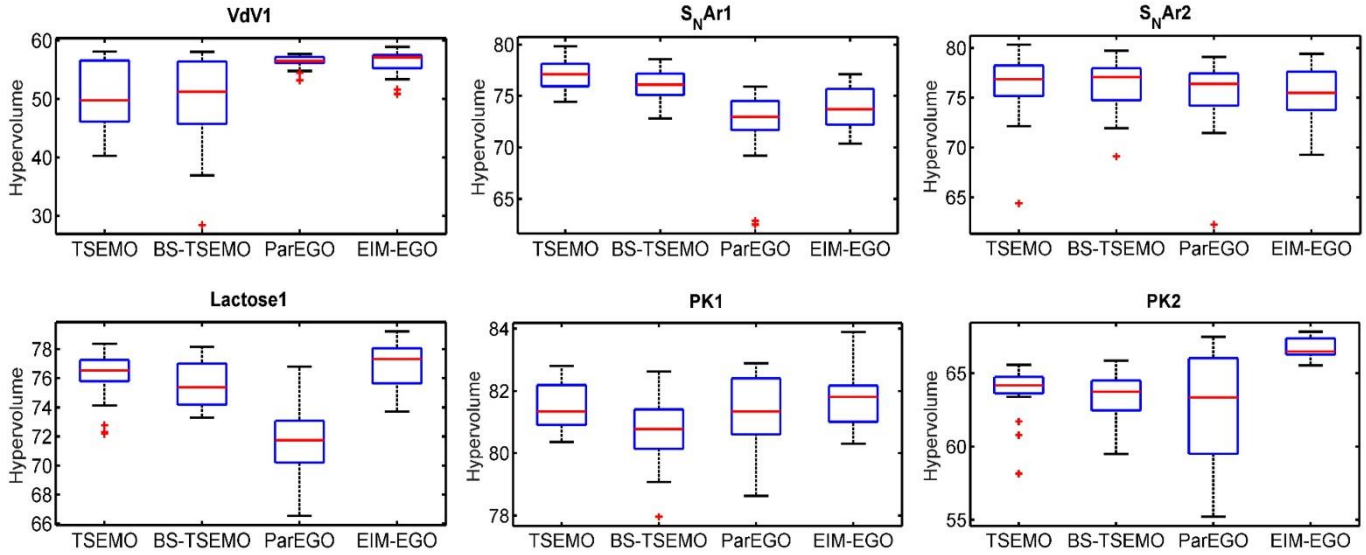

Figure 3. Boxplots of the optimisation results after 60 function evaluations using hypervolume as a performance indicator (20 runs per algorithm). The NSGA-II results were omitted for clarity. + = outlier (more than 1.5× the interquartile range away from the upper or lower quartile).

To validate the use of the MOOA reaction simulator for identifying the best algorithm to use for reaction optimisation, we compared the performance of TSEMO and EIM-EGO on a three-objective experimental optimisation, with the expectation that TSEMO would outperform EIM-EGO.

Experimental Comparison

Methyl phenyl sulfide can be oxidised by hydrogen peroxide to the sulfoxide and subsequently to the sulfone in continuous flow (Figure 4). Often the sulfone is an undesired side product of this reaction, for example, the oxidation of sulfur mustard forms the significantly less hazardous sulfoxide, but this can be further oxidised to form the highly toxic sulfone.[34] Often the solution for selective oxidation is to use a catalyst, or to conduct the reaction in a polar protic solvent.[35–37] Herein, we focus on the optimisation of methyl phenyl sulfide conversion, selectivity and space-time-yield (STY) using temperature, residence time and equivalents of H_2O_2 compared to the sulfide as reaction variables.

The experimental study was performed on a flow platform using an inline UV detector for residence time studies and a GCMS fitted with an automated liquid samplier for determining conversion and selectivity. (Figure 4). In detail, reagents were pumped using JASCO PU2080 and PU2085 dual piston HPLC pumps and pump streams were mixed using Swagelok SS-100-3 tee-pieces. A tubular reactor of desired volume was fitted to an aluminium cylinder heated by heating cartridges controlled by a Eurotherm 980. Sampling was achieved using a Shimadzu Flowcell, from which the AOC 6000 auto sampler extracts a volume of 0.5 µL. The reactor was maintained under the desired fixed back pressure using an Upchurch Scientific psi back pressure regulator. Polyflon PTFE tubing (1/16" OD, 1/32" ID) was used throughout the reactor. Quantitative analysis was performed on a Shimadzu GCMS-QP2010 SE instrument fitted with a SH-Rxi-5Sil MS column (30 m length, 2.5 mm ID and 2.5 µm film thickness), α,α,α-trifluorotoluene (3FB) was included as an internal standard.

Ideal plug flow could be assumed as characterisation via residence time distribution experiments found Bodenstein numbers to be above 100 (SI section 3.3).[38] Additionally, a time equivalent of 1.5 reactor volumes was allowed to elapse before sampling.[39] Flow rates, reaction temperature and sampling were monitored by the MATLAB interface, which used the calibrated mass spectrum peak areas for analysis, with either the TSEMO or EIM-EGO algorithm to calculate the next reaction setting. For EIM-EGO the Euclidean distance-based function was chosen as the acquisition function to conform to the prior simulations. More experimental details can be found in the SI.

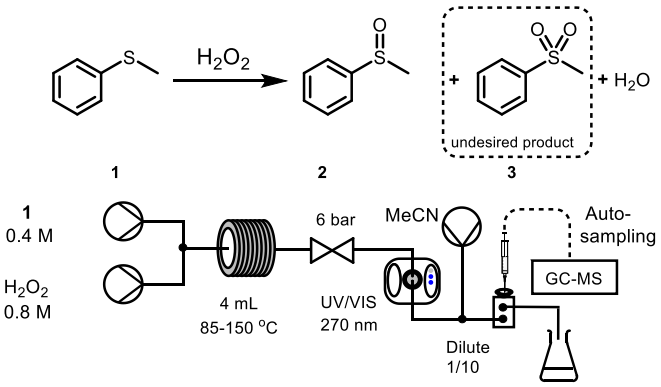

Figure 4. Reaction and experimental scheme for methyl phenyl sulfide, **1**, automated oxidation to methyl phenyl sulfoxide, **2**, and successively methyl phenyl sulfone, **3**. The dilution is added allowing direct sampling to the GCMS. Samples are taken by an automated liquid sampler (Shimadzu AOC-6000).

Both algorithms were initialised by a test set of 15 experiments created by the MATLAB LHC function with 1000 iterations for an optimal distribution of experiments. Additionally, reproducibility was tested by repeating one experimental condition six times during the optimisation. Two of those

Please do not adjust margins

experiments were added to the initial test set. Conversion, selectivity and STY could be reproduced with less than 5% deviation (Figure S10 SI).
TSEMO and EIM-EGO both ran 35 single experiment iterations. The optimisation experimental results are shown in Figure 5.

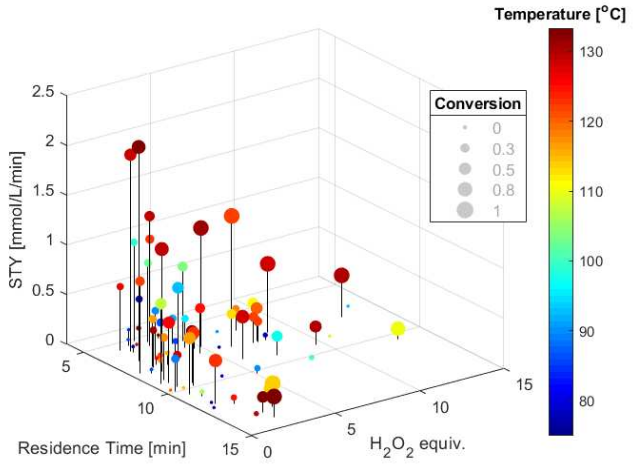

Figure 5. Overview of TSEMO and EIMEGO optimisations. Results show dependence of STY, of the formation of methyl phenyl sulfoxide and the conversion of methyl phenyl sulfide on temperature, residence time and hydrogen peroxide equiv.

The optimisations found an optimal region at temperatures above 115 °C, residence times around 6.5 minutes and H_2O_2 equivalents of around 2.5. Not plotted is the selectivity because in 2/3 of the optimisation experiments selectivities of over 90 % were observed. A full overview of all parameter and objective relations can be found in the SI. Our results agree with literature showing equivalent ratios of 2.5 as optimal for highly selective formation of sulfoxide at medium temperatures.[37,40] An explanation can be potentially found in mechanistic studies describing the O-O bond breaking as rate limiting and often stabilized by water.[41,42]

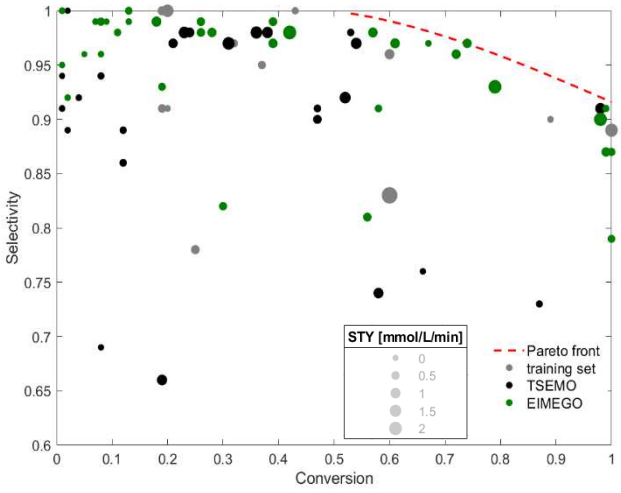

Figure 6. Approximated Pareto front and experimental data near it. Conversion, Selectivity and STY of **2** were optimised against each other.

Hence, the equivalence could be a summation of water and hydrogen peroxide molarity in the solution (0.8 M H_2O_2 and the usually neglected ~0.3 M water).
Based on both algorithms a Pareto front for the 2-dimensional correlation between conversion and selectivity can be easily visually depicted (Figure 6). For the approximation in 3-dimensional space compare with Figure S14 in the SI.
The results detail correlation between all three optimisation goals and optimisation trade-offs. Therefore, reducing experimental load and speculation compared to traditional experimental approaches. Both algorithms found similar results regarding optima, but their speed and approach in covering the most hypervolume varied.
TSEMO optimised for the maximum hypervolume faster and more effectively (Figure 7). Looking at the iteration steps following the initial test set, (SI, section 4) a more risk-taking approach can be distinguished. While EIMEGO initially stayed close to already found optima at average conditions, TSEMO tested more extreme points closer to the boundaries faster.

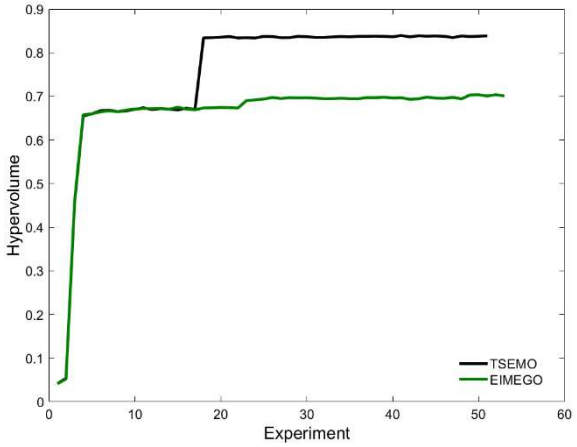

Figure 7. Change of hypervolume for TSEMO and EIMEGO throughout the optimisations.

These experimental results support the theoretically found results of TSEMO being superior to EIMEGO for a 3-objective problem.

Conclusions

The efficiency of self-optimisation can be directly related to the choice of optimisation algorithm. Therefore, it is important to keep self-optimising systems up to date with the latest advances in computer science. A method to compare the performance of multi-objective optimisation algorithms for the self-optimisation of chemical reactions was developed using a kinetic-based reaction simulator. The simulator included six chemistry-inspired test problems, each representing different types of reactions and possessing Pareto fronts with different properties. This approach was used to compare five different MOOAs, and identify which algorithm was best for each type of test problem. All Bayesian optimisation algorithms succeeded in finding an optimum quickly compared to NSGA-II, which can be considered less favourable for expensive to evaluate problems. EIM-EGO was in all cases the fastest, apart from for 3 objective

Please do not adjust margins

optimisations, for which TSEMO was better. The oxidation of methyl phenyl sulfide by hydrogen peroxide was therefore optimised for conversion, selectivity and space time yield of meythl phenyl sulfoxide using the two different algorithms. Both algorithms deepened the understanding of this reaction, revealing correlations and limitations of the reaction parameters with respect to the optimisation goals. The simulations and experimental case study both favoured TSEMO over EIM-EGO in succeeding in finding an optima in the shortest time. We propose this is due to a small difference in the optimization of exploration over exploitation. The outcome was in agreement with predictions made from prior simulations, thus demonstrating the usefulness of the multi-objective reaction simulator. Future work will focus on expanding the capabilities of the reaction simulator to encompass a wider range of more complex chemical examples (catalytic cycles etc.) The question, which algorithm to choose cannot be answered universally, the presented approach can however predetermine *in silico the* most suitable one for a specific challenge. Thus increasing the efficiency of subsequent experimental optimisations.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

PM thanks Defence Science and Technology Laboratory and University of Leeds for funding. ADC thanks UCB, AstraZeneca and University of Leeds for funding. SVL thanks the ACS Authur Cope Award for funding.

References

1 C. W. Coley, D. A. Thomas, J. A. M. Lummiss, J. N. Jaworski, C. P. Breen, V. Schultz, T. Hart, J. S. Fishman, L. Rogers, H. Gao, R. W. Hicklin, P. P. Plehiers, J. Byington, J. S. Piotti, W. H. Green, A. J. Hart, T. F. Jamison and K. F. Jensen, *Science*, 2019, **365**, 1–9.
2 J. M. Granda, L. Donina, V. Dragone, D.-L. Long and L. Cronin, *Nature*, 2018, **559**, 377–381.
3 S. Steiner, J. Wolf, S. Glatzel, A. Andreou, J. M. Granda, G. Keenan, T. Hinkley, G. Aragon-Camarasa, P. J. Kitson, D. Angelone and L. Cronin, *Science*, 2019, **363**, 1–8.
4 S. Chatterjee, M. Guidi, P. H. Seeberger and K. Gilmore, *Nature*, 2020, **579**, 379–384.
5 D. Perera, J. W. Tucker, S. Brahmbhatt, C. J. Helal, A. Chong, W. Farrell, P. Richardson and N. W. Sach, *Science*, 2018, **359**, 429–434.
6 C. Mateos, M. J. Nieves-Remacha and J. A. Rincón, *React. Chem. Eng.*, 2019, **4**, 1536–1544.
7 B. J. Shields, J. Stevens, J. Li, M. Parasram, F. Damani, J. I. M. Alvarado, J. M. Janey, R. P. Adams and A. G. Doyle, *Nature*, 2021, **590**, 89–96.
8 A. M. Schweidtmann, A. D. Clayton, N. Holmes, E. Bradford, R. A. Bourne and A. A. Lapkin, *Chem. Eng. J.*, 2018, **352**, 277–282.
9 A. D. Clayton, A. M. Schweidtmann, G. Clemens, J. A. Manson, C. J. Taylor, C. G. Niño, T. W. Chamberlain, N. Kapur, A. J. Blacker, A. A. Lapkin and R. A. Bourne, *Chem. Eng. J.*, 2020, **384**, 123340.
10 J. A. Manson, T. W. Chamberlain and R. A. Bourne, *J. Glob. Optim.*, 2021, **80**, 865–886.
11 A. D. Clayton, J. A. Manson, C. J. Taylor, T. W. Chamberlain, B. A. Taylor, G. Clemens and R. A. Bourne, *React. Chem. Eng.*, 2019, **4**, 1545–1554.
12 K. C. Felton, J. G. Rittig and A. A. Lapkin, *Chemistry–Methods*, 2021, **1**, 116–122.
13 A. S. Surur, L. Schulig and A. Link, *Arch. Pharm. (Weinheim)*, 2019, **352**, 1800248.
14 R. V. Kupwade, *J. Chem. Rev.*, 2019, **1**, 99–113.
15 S. L. Bartelt-Hunt, D. R. U. Knappe and M. A. Barlaz, *Crit. Rev. Environ. Sci. Technol.*, 2008, **38**, 112–136.
16 H. Wang, G. W. Wagner, A. X. Lu, D. L. Nguyen, J. H. Buchanan, P. M. McNutt and C. J. Karwacki, *ACS Appl. Mater. Interfaces*, 2018, **10**, 18771–18777.
17 M. Jereb, *Green Chem.*, 2012, **14**, 3047–3052.
18 S. V. Ley, D. E. Fitzpatrick, R. J. Ingham and R. M. Myers, *Angew. Chem. Int. Ed.*, 2015, **54**, 3449–3464.
19 J. Vojtesek and P. Dostal, Vysna Boca, 2010, p. 10.
20 J. W. Lee, Z. Horváth, A. G. O'Brien, P. H. Seeberger and A. Seidel-Morgenstern, *Chem. Eng. J.*, 2014, **251**, 355–370.
21 S. A. Hashemi and F. Z. Ashtiani, *Food Bioprod. Process.*, 2010, **88**, 181–187.
22 J. S. Moore and K. F. Jensen, *Angew. Chem. Int. Ed.*, 2014, **53**, 470–473.
23 S. Mirjalili and A. Lewis, *Inf. Sci.*, 2015, **300**, 158–192.
24 C. A. Hone, Thesis, University of Leeds, 2016.
25 R. M. Everson, J. E. Fieldsend and S. Singh, Springer, London, UK, 2002, pp. 343–354.
26 K. Bringmann, T. Friedrich, C. Igel and T. Voß, *Artif. Intell.*, 2013, **204**, 22–29.
27 J. Knowles, *IEEE Trans. Evol. Comput.*, 2006, **10**, 50–66.
28 E. Bradford, A. M. Schweidtmann and A. Lapkin, *J. Glob. Optim.*, 2018, **71**, 407–438.
29 D. Zhan, Y. Cheng and J. Liu, *IEEE Trans. Evol. Comput.*, 2017, **21**, 956–975.
30 K. Deb, A. Pratap, S. Agarwal and T. Meyarivan, *Evol. Comput. IEEE Trans. On*, 2002, **6**, 182–197.
31 Y. Tian, R. Cheng, X. Zhang and Y. Jin, *IEEE Comput. Intell. Mag.*, 2017, **12**, 73–87.
32 Y. Amar, A. M. Schweidtmann, P. Deutsch, L. Cao and A. Lapkin, *Chem. Sci.*, 2019, **10**, 6697–6706.
33 V. R. Joseph and Y. Hung, *Stat. Sin.*, 2008, **18**, 171–186.
34 M. A. Stahlmann, C. Golumbic, W. Stein and J. S. Fruiton, *J. Org. Chem.*, 1946, **11**, 719–735.
35 S. Doherty, J. G. Knight, M. A. Carroll, J. R. Ellison, S. J. Hobson, S. Stevens, C. Hardacre and P. Goodrich, *Green Chem.*, 2015, **17**, 1559–1571.
36 C. Ren, R. Fang, X. Yu and S. Wang, *Tetrahedron Lett.*, 2018, **59**, 982–986.
37 K. Sato, M. Hyodo, M. Aoki, X.-Q. Zheng and R. Noyori, *Tetrahedron*, 2001, **57**, 2469–2476.
38 P. Müller, A. Aguirre, T. Ljungdahl and J. van der Schaaf, *J. Adv. Manuf. Process.*, 2020, **2**, 1–9.
39 G. Vernet, M.-S. Salehi, P. Lopatka, S. K. Wilkinson, S. K. Bermingham, R. Munday, A. O'Kearney-McMullan, K. Leslie, C. A. Hone and C. Oliver Kappe, *Chem. Eng. J.*, 2021, **416**, 129045.
40 S. Doherty, J. G. Knight, M. A. Carroll, A. R. Clemmet, J. R. Ellison, T. Backhouse, N. Holmes, L. A. Thompson and R. A. Bourne, *RSC Adv.*, 2016, **6**, 73118–73131.
41 J.-W. Chu and B. L. Trout, *J. Am. Chem. Soc.*, 2004, **126**, 900–908.

Please do not adjust margins

42 R. Wang, L. Lu, J. Chen, J. Kou, H. Jin and L. Guo, *J. Mol. Liq.*, 2021, **330**, 115654.

Please do not adjust margins