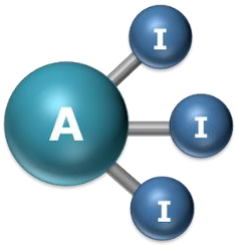

Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

Optimising Ag/Au Alloyed Nanoparticle Catalysts in Continuous Flow; Discrete vs. Continuous Variable Optimisation

Project Report

[Project Dates: 23/08/2021 – 01/10/2021

TWC group, iPRD lab, School of Chemistry, University of Leeds

Project Student: Louis Greenhalgh, Student at University of Leeds

Supervised by: Dr. Thomas Chamberlain, Associate Professor at University of Leeds

Report Date: 13/10/2021

AI3SD-Intern-Series:Report-4_Greenhalgh

Optimising Ag/Au Alloyed Nanoparticle Catalysts in Continuous Flow; Discrete vs. Continuous Variable Optimisation
AI3SD- Intern-Series:Report-4_Greenhalgh
Report Date: 13/10/2021
DOI: 10.5258/SOTON/AI3SD0146
Published by University of Southampton

Network: Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery
This Network+ is EPSRC Funded under Grant No: EP/S000356/1
Principal Investigator: *Professor Jeremy Frey*
Co-Investigator: *Professor Mahesan Niranjan*
Network+ Coordinator: *Dr Samantha Kanza*

Table of Contents

1. Project Details **1**

2. Project Team **1**

2.1 Project Student **1**

2.2 Supervisor **1**

2.3 Researchers & Collaborators **1**

3. Lay Summary **1**

4. Aims, Objectives and Introduction **2**

5. Methodology **5**

6. Results **100**

7. Conclusions & Future Work **18**

8. Outputs, Data & Software Links **19**

9. References and Bibliography **19**

1. Project Details

Title	Optimising Ag/Au Alloyed Nanoparticle Catalysts in Continuous Flow; Discrete vs. Continuous Variable Optimisation
Project Reference	AI3SD-FundingCall3_004
Supervisor Institution	University of Leeds
Project Dates	23/08/2021 – 01/10/2021
Website	https://chamberlainresearchgroup.com/
Keywords	Nanoparticles, continuous flow, self-optimisation, BayesOpt

2. Project Team

2.1 Project Student

Name and title		**Louis Greenhalgh**
Employer name/University Department Name		**University of Leeds School of Chemistry**
Work Email: **cm18lg@leeds.ac.uk**	**Website Link (if available)**	**n/a**

2.2 Project Supervisor

Name and title:		**Dr. Thomas Chamberlain**
Employer name/University Department Name		**University of Leeds School of Chemistry [cit.]**
Work Email: **t.w.chamberlain@leeds.ac.uk**	**Website Link (if available)**	**https://chamberlainresearchgroup.com/**

2.3 Researchers & Collaborators

Dr. Thomas Chamberlain (supervisor), Brendon Hall (University of Leeds), Jamie Manson, Luke Power (University of Leeds)

3. Lay Summary

Incorporating artificially intelligent optimisation algorithms into chemistry experiments can allow scientists to reduce the number of man-hours and materials used, by enabling the smart selection of which experiments to run in order to more efficiently reach experiment outcomes. This has been combined with continuous-flow reactor technology, which has great potential to automate scientific discovery. There is also the potential to obtain understanding for more efficient scale-up towards industry, where small materials savings have large economics of scale. In this project, AI optimisation was used to determine the benefits of continuous over discrete variable optimisation of the ratio of gold and silver content in nanoparticles used as catalysts in continuous-flow chemical synthesis.

4. Aims, Objectives, and Introduction

Aims and Objectives:

The aim of this project is to provide evidence that continuous-variable optimisation using a previously devised continuous-optimisation algorithm, is more efficient at reaching a global optimum than discrete optimisation, which is exemplified using a novel discrete Bayesian optimisation algorithm.
The subject for this investigation was optimising the ratio of gold:silver in a solution of catalytic colloidal nanoparticles, alongside concentration of reducing agent ($NaBH_3$) and residence time, with the aim of maximising conversion of nitrophenol to phenylamine in an example reaction in continuous flow. Previous attempts at optimisation of the Au:Ag ratio were done continuously with much success.
This project aims to confirm the benefits of continuous optimisation over discrete optimisation, by using a mixed continuous & discrete Bayesian optimisation algorithm and comparing the results. As a secondary goal, the number of discrete levels explored by the algorithm would be varied, implementing linearly as well as exponentially spaced discrete levels. A tertiary goal would be to investigate the differences in exploration of experimental space between experiments with varied numbers of discrete levels, and compare between them and the former successful continuous approach.

Introduction:

The application of the ideas resulting from such an investigation fits into a more general problem in both scientific discovery and industry, namely that of cost. The implications of reducing process cost at the large scale of chemical industry are obvious, while the impact on research is less so.
Superficially, scientific research and discovery may seem less concerned with economics, however the ability of university research groups and research institutes to fund their research is often greatly impacted by cost/benefit analysis. Research that has clearer economic potential in industry is far more likely to receive a steady supply of funding than theoretical research with less clear commercial application. This is a problem that deserves attention, since it is not always obvious what early-stage, fundamental discoveries can lead to world-changing technologies in the future.

An answer to this problem involves expediting scientific discovery using technology to reduce the costs associated with scientific research. This involves looking at the fundamentals of science, the scientific process (i.e., *how* experiments are performed and *how* science is actually conducted), and of the role of the scientist; all of which is the subject of metascience.
Optimisation through machine-learning tips the balance and alters the role of the scientist away from the steering wheel in low-level experimental design towards higher level automation systems, simultaneously reducing human input, material cost and environmental impact of both experimentation and scaled-up chemical process.

The scientific method relies upon the principal of the dependant and independent variable (also called the "response variable") to prove or disprove hypotheses. By holding steady the other variables affecting the conditions of a system, the dependent variables, in such a way that changing one variable, the independent variable, will provide insight to the dependence of the system's conditions *on that independent variable.* This provides the basis of experimentation, and refines what could be a chaotic mess of trial and error into a

systematic protocol aimed at isolating causes from effects in the pursuit of *validity, reliability, and replicability* of results.

The role of the scientist is to decide the independent variable in line with their hypothesis, isolate the system from the dependent variables that influence it, and finally to decide the area of experimental space they wish to explore with their experiment. Experimental space is rather like a sample space in statistics. A sample space is the set of all possible outcomes of a scenario, much like experimental space is the set of all possible *results* within a single experimental design space, which has been designed in line with the earlier described principals. Optimising the exploration efficiency of experimental space has already been a topic in scientific discourse for decades [1][2], termed "design of experiments" or "DOE".

DOE involves establishment of a response surface, which enables a graphical representation of the dependence of the system on multiple response variables. A DOE visualisation across experimental space demonstrates the difficulties associated with chemical system optimisation, and why there exists no single method to optimise a set of experiments in chemistry [6].

Traditionally, in multivariable systems an analogue "one variable at a time" (OVAT), approach would be used to optimise a chemical system. This is extremely time consuming, and involves a trade-off between the amount of experimental space explored and the number of experiments that are to be performed, i.e., where researchers must decide where to allocate their limited time and access to resources.
This approach also rarely leads to the discovery of the most optimum reaction conditions in a chemical system, as shown in figure 1.

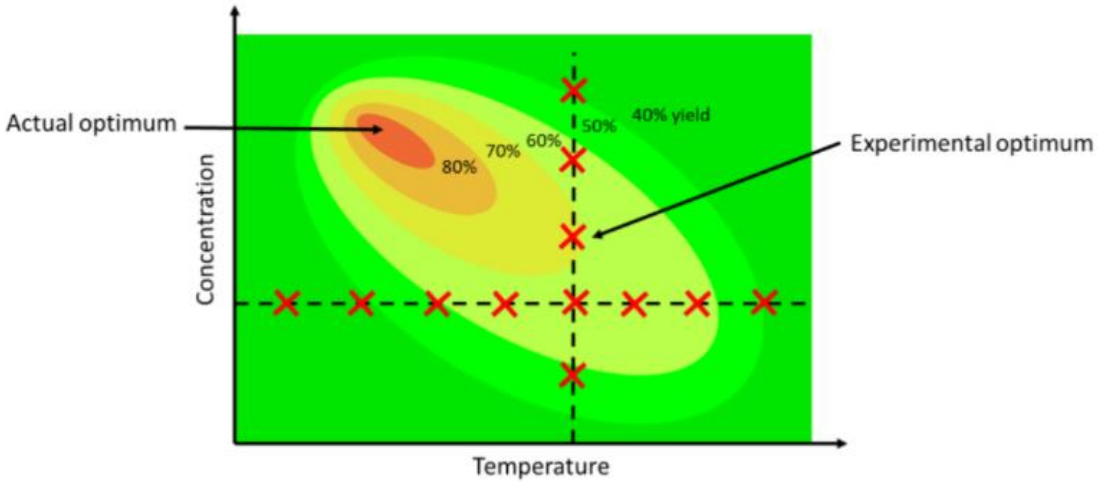

Figure 1: A response surface that exemplifies how an OVAT approach to multivariable optimisation easily misses the true global optimum, where crosses indicate individual experiments in a chosen DOE. A higher yield in a prior experiment at lower concentration changes the point of focus such that the area of experimental space containing the true global optimum would be missed on subsequent experiments at higher concentration.

Optimisation of a multivariable chemical system may involve a trade-off between one variable and another, meaning that a single solution that simultaneously optimises every variable may not exist. This is where the objective functions of the variables are said to conflict. In this case, one variable cannot be improved without worsening another, which results in what is known as a "Pareto front" of "non-dominated" solutions [3].

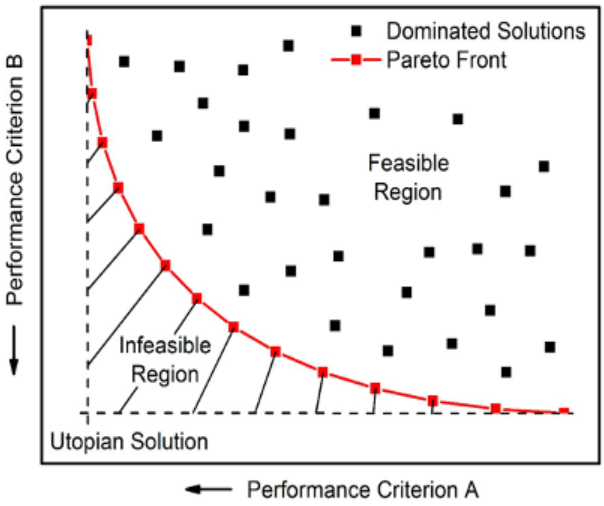

Figure 2: A Pareto front of non-dominated solutions from performance criteria A & B, bordered by dominated solutions in a feasible region, and an infeasible region. This is the border between where one criterion's objective function cannot be minimised further without raising the other, and infeasibility.

Since there are potentially an infinite number of solutions along the Pareto front, the scientist must decide on a solution that aligns with the aims of their experimentation. For example, many synthetic chemical experiments in research aim to increase yield or conversion of a starting material. This could be at the expense of residence time in a continuous flow reactor; an effect which could be amplified at industrial scale determining rate of production in a chemical plant. This was shown in 2013, where Ganesan *et al.* performed a multi-objective optimisation of carbon dioxide reforming and partial oxidation of methane, simultaneously optimising methane conversion, carbon monoxide selectivity and hydrogen to carbon monoxide ratio [4].

Optimisation through machine learning uses the data produced through experiment as training data for a machine learning algorithm, which is used on the fly to generate a surrogate model of the system. A surrogate model can be thought of as the bridge between the algorithm's view of the chaotic, noisy, real-life chemical system, the "objective function", and the set of experimental results arising from that system. It is a simplified version of the real-life system required to perform the statistical calculations necessary within an optimisation algorithm. This enables the minimisation of the objective function along the Pareto front through iterated improvement of the surrogate model in Bayesian optimisation.

The surrogate model is a Gaussian process where the distribution of possible results is dealt with as an infinite array of Gaussian functions, represented by a mean and covariance function. This is used in combination with an acquisition function to automate the process of experiment selection, which is tuned using hyperparameters that determine the balance between the tendency of the acquisition function to explore new experimental space (and avoid getting stuck in local optima) vs. minimising the number of iterations it takes to reach a global optimum (a common measure of optimisation algorithm efficiency).

Bayesian optimisation updates the surrogate model of the objective function on each iteration after each experiment, using the new data to inform the selection of the next

experiment, with the ultimate aim of minimising the objective function to reach a global optimum. This is done by producing a surrogate model where the distribution is closer to that of the true objective function, where variance (representing uncertainty) will be greater further away from the training data.
It is termed "Bayesian" because priors (the previous surrogate model) are updated on the input of new data, to produce the updated surrogate model. The algorithm will run either until a certain optimum or iteration limit is reached, either in a simulation or a set reasonable number of experiments in real life.

5. Methodology

In this project the response variable is conversion percentage of nitrophenol to phenylamine, measured using on-line UV-Vis spectrometry. The 3 dependent variables are: the continuous residence time of the reaction mixture in the reactor system; discrete levels of the Au:Ag composition ratio of the alloyed nanoparticle catalyst; and the continuous ratio of nanoparticle catalyst:reducing agent, $NaBH_3$. This is termed "mixed variable multi-objective optimisation" or "MVMOO", where optimisation is performed between multiple variables of differing type.[5]

The Bayesian MVMOO algorithm used in this project was created by Jamie Manson using MATLAB [5], and utilised a surrogate model generated from a sample of training data stemming from a previous continuous-flow reactor system built by Brendon Hall. He used SNOBFIT (another optimisation algorithm) with real-life experiments in the reactor to optimise the continuous ratio of gold:silver in alloyed nanoparticles produced using the Turkevitch protocol[7]. B. Hall chose to explore a resolution of 1% Au:Ag above that granted by the minimum pump steps in his continuous flow system, which was theoretically 0.005%.

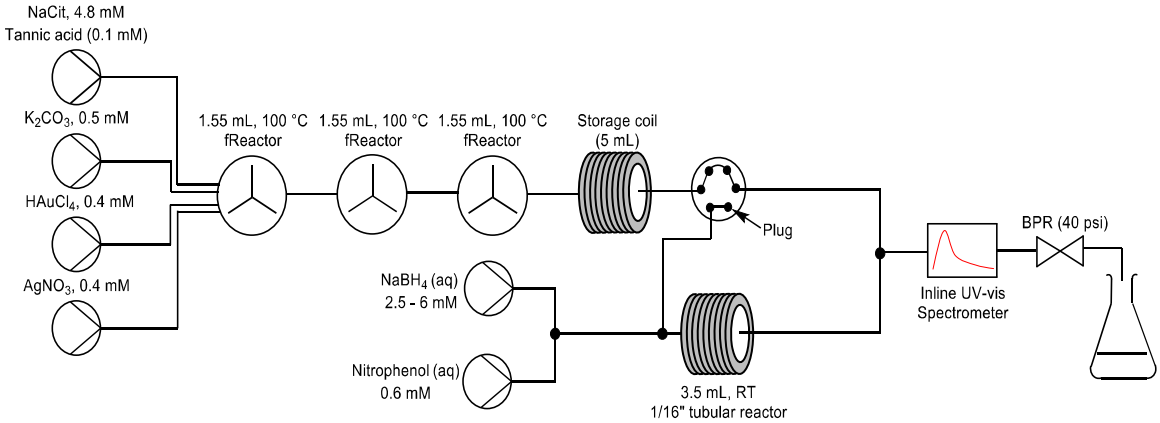

Figure 3: Self-optimising continuous-flow reactor schematic [6]

Table 1: Reactor data from the continuous optimisation of Au:Ag in a self-optimising continuous-flow reactor [6]

Au:Ag ratio	AuAgNP:$NaBH_4$ ratio	Residence time (min)	Conversion (%)
0.814724	0.228864	0.922886	43.0696
0.031833	0.152442	1.443623	24.286
0	0.25	0.53	27.9091
0.93	0.24	1.5	39.2271
0.3	0.15	0.98	20.5675
0.52	0.16	0.55	28.1981
0.96	0.17	0.52	38.7054
0.5	0.23	1.46	30.1455
0	0.27	0.98	28.434
0.54	0.11	1.17	29.4197
1	0.3	0.5	48.2176
1	0.29	0.96	57.0241
0.23	0.23	1.22	22.9256
1	0.3	1.1	60.2201
1	0.3	1	53.5624
0.62	0.2	1.07	37.5694
1	0.28	1.5	56.5323
1	0.21	1.17	42.626
0.23	0.13	1.22	24.4031
1	0.27	1.2	56.7796
1	0.2	0.89	39.0458
0.12	0.27	1.36	24.9779
1	0.21	0.7	42.1776
0.981472	0.196852	1.249751	26.66667
0.982582	0.141601	1.991816	49.5
0.93	0.08	0.71	19.91667
0.94	0.03	1.62	17.91667
0.91	0.04	1	15.16667
0.99	0.28	1.75	60.25
0.91	0.3	0.89	29.75
0.95	0.3	1.48	52.66667
0.92	0.01	1.38	13.33333
0.93	0.02	1.83	19.58333
1	0.01	0.89	10.33333
1	0.3	2	54.91667
0.94	0.16	1.44	36.41667
0.99	0.01	1.85	15.41667
1	0.3	1.76	52
0.93	0.27	1.23	37.83333
0.97	0.3	1.5	64
0.96	0.3	1.87	62.33333
0.92	0.2	1.21	31.66667
0.96	0.3	1.21	51.75
0.96	0.27	1.7	59.91667
0.96	0.25	0.79	44.41667
0.98	0.3	1.79	64.25
0.97	0.16	1.44	46.08333

0.98	0.3	2	66.41667
0.96	0.15	1.82	46.58333
0.96	0.08	1.16	27.08333
0.97	0.3	1.99	66.58333
0.98	0.22	1.79	55.58333
0.99	0.2	1.53	45.66667
0.98456	0.030362	1.7153	23.0989
0.92958	0.23741	1.0063	34.1548
0.94606	0.26613	1.4542	49.9892
0.96447	0.18421	0.72304	41.7431
0.99377	0.10759	1.155	29.4664
0.9079	0.16003	1.3273	29.45
0.9242	0.059235	1.891	21.4767
0.99324	0.3	1.9133	58.0437
0.93738	0.3	1.7192	49.142
0.97086	0.3	2	72.9477
0.97178	0.3	0.80842	56.0847
0.96561	0.3	2	71.9256
0.97164	0.26905	2	68.9019
0.97227	0.3	2	71.9945
0.96985	0.29998	2	72.6214

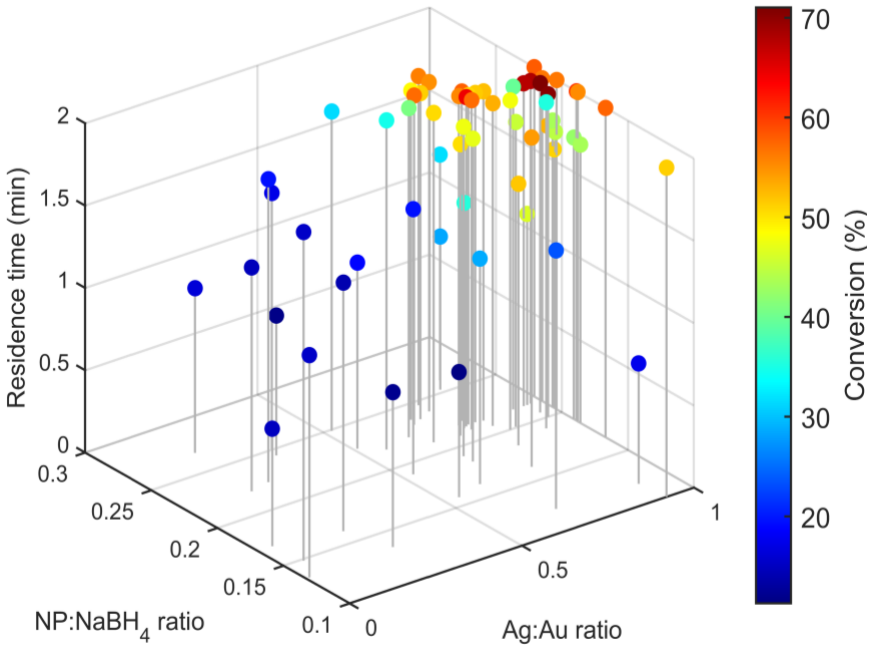

Figure 4: B. Hall's continuously optimised data from real-life experiment on the cont. flow reactor [6]

Brendon received a surprising optimum conversion percentage with 97% gold nanoparticles. This demonstrated the power of continuous optimisation since the resolution of most OVAT approaches would prohibit access to the experimental space containing the global optimum without inordinate effort and cost.

The Bayesian MVMOO algorithm used in this project was set with varying numbers of allowed discrete levels for the AuAg ratio of the nanoparticles, essentially minicking the experiment in which batches of nanoparticle catalysts were prepared and made available for testing, e.g. from wells in an autosampler. The experiments were run *in silico* and run for 50 iterations for 5 times per setting. The bounds of the simulation were set to the same as those used by Brendon in his reactor. This was done with the prior knowledge of the response surface and global optimum found by Brendon, therefore with the objective of ascertaining the difficulty of finding that same optimum through discretely varying the Au:Ag ratio.
A limitation of this approach was that the program in MATLAB could not handle an infinite number of discrete levels, since this exponentially increased the computational workload of the system. It was found that beyond 20 discrete levels, the algorithm would still produce results, but with very few algorithm solver runs returning positive exit flags, indicating that the results obtained were unreliable. Past 50 discrete levels, the program would break down, crash the computer, and could not produce meaningful results.

The number of discrete levels were varied between 5, 11, 15 and 20 levels.
This was split into 2 methods:

Method 1: N discrete linearly spaced levels between 0 and 1
Method 2: N Linearly spaced values of x between 0 and 2.3 were substituted into equation 1 to produced N exponentially spaced levels. (e.g. 5 linearly spaced values between 0 and 2.3 plugged into eqn. 1 produced the levels).

$$\boldsymbol{y(x) = -10^{-x} + 1} \qquad \text{(Equation 1)}$$

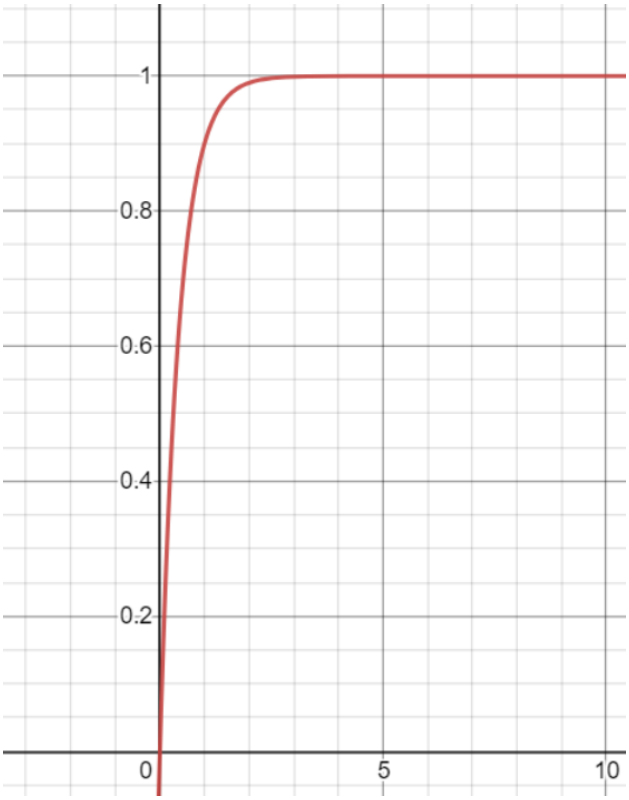

Figure 5: Graph of function shown in equation 1 used to produce exponentially spaced discrete levels, chosen between x = 0 and 2.3, to receive y(x) = 0.005

The aim of this approach was to scan the experimental space accessed by the simulation with method 1, increasing from 5 up to 20 discrete levels. Method 2 was an extension designed to explore the effect of tuned exponential increase in the percentage of gold in the nanoparticles, between 0% and 99.5% gold. This was achieved with a script that populated an array with varied numbers of discrete values for the levels with linearly spaced substitutions of x along the function.

The justification for narrowing the experimental space explored in this way was two-fold: since the system had previously been found to be more heavily reliant on the presence of gold in the nanoparticles, it made sense to produce a system that skewed its focus of exploration towards the higher ratios of gold:silver.
Another reason was more general, since there are many situations in chemical systems where more = better. In chemical kinetics, reaction rate is proportional to concentration, and there exist many catalytic systems in which higher amounts of catalyst will result in a faster reaction. This means such an approach of tuned exponentially increasing discrete levels could show application in future experimental design if it turns out to be more efficient at exploring experimental space in this higher region.

Finally, as another extension to the project, the exploration of experimental space was investigated using a custom-built script based on ideas from Luke Power of an exploration-factor. This value was calculated based on the average shortest distance of each explored point in experimental space from every other point, over the total experimental space explored:

$$\boldsymbol{Exploration\ factor}, \boldsymbol{E} = \frac{\sum \boldsymbol{shortest\ distance}}{\boldsymbol{Area\ explored}} \qquad \text{(Equation 2)}$$

A higher average distance equals a higher average exploration across the space.

The results from the simulated experiments were plotted on a 3D graph, with colour in MATLAB to compare the maximum conversion percentage achieved on each run. Cumulative conversion percentage was also plotted as a measure of optimisation efficiency, to determine how quickly (i.e. after how many experiments) the optimisation achieved the optimum.

6. Results

N Linearly spaced levels:

5 levels of Au:Ag Ratio

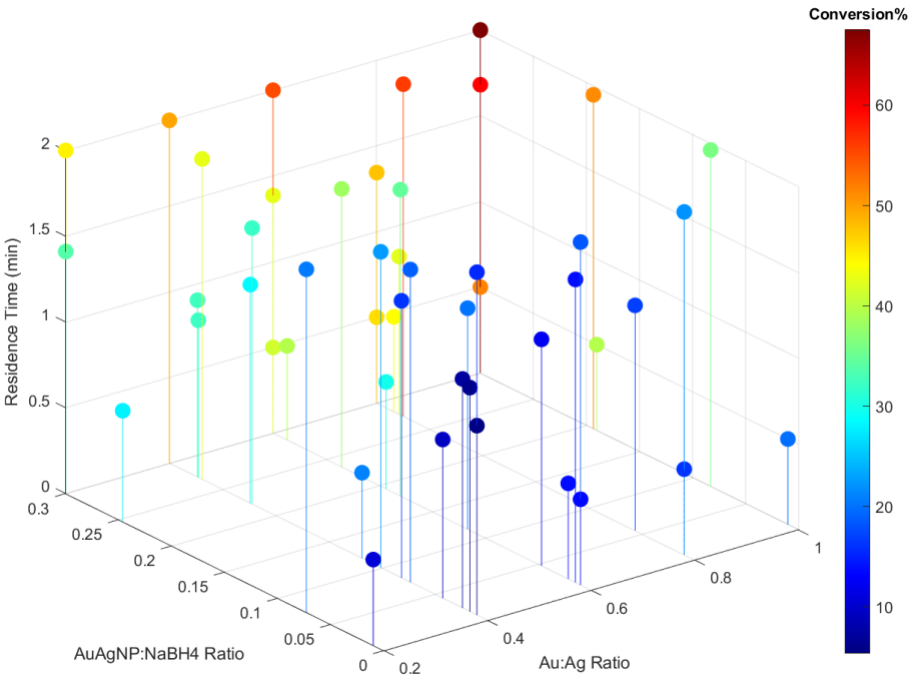

Figure 6: Response surface resulting from optimisation involving 5 linearly spaced levels of Au:Ag ratio in a nanoparticle catalyst.

11 levels of Au:Ag Ratio

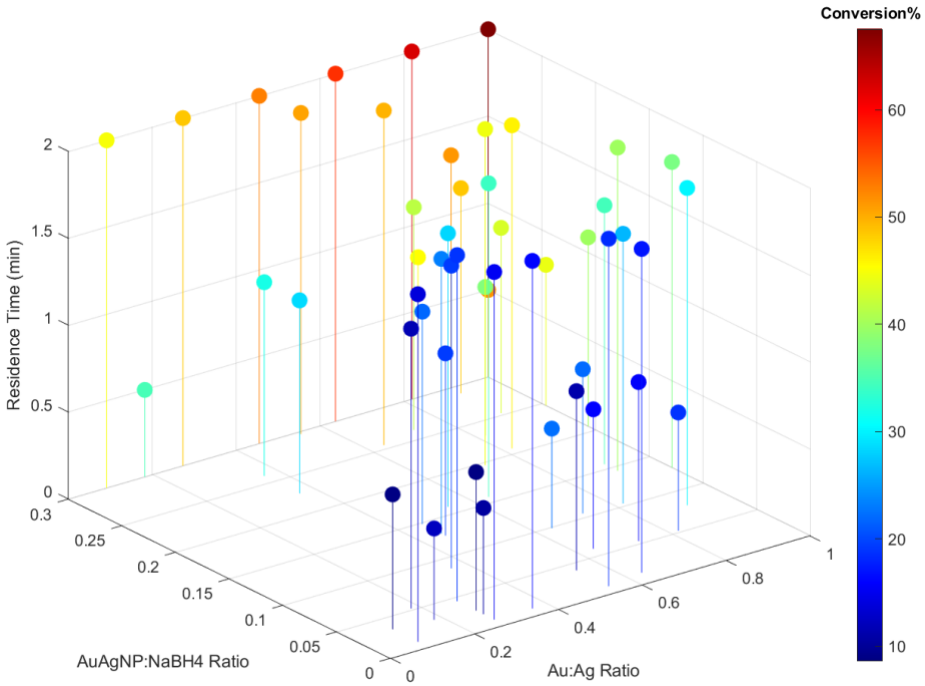

Figure 7: Response surface resulting from optimisation involving 11 linearly spaced levels of Au:Ag ratio in a nanoparticle catalyst.

15 levels of Au:Ag Ratio

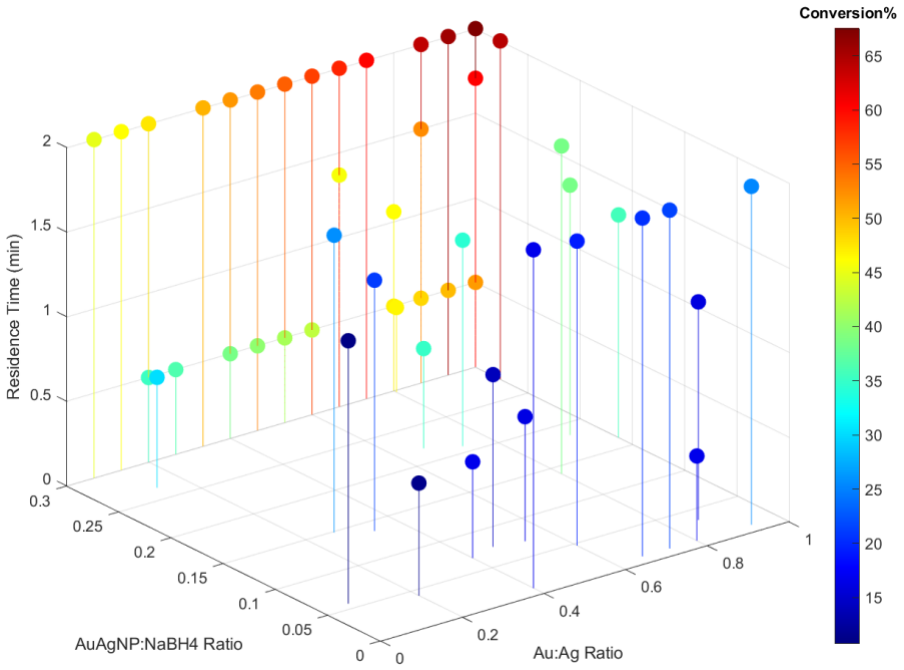

Figure 8: Response surface resulting from optimisation involving 15 linearly spaced levels of Au:Ag ratio in a nanoparticle catalyst.

21 levels of Au:Ag Ratio

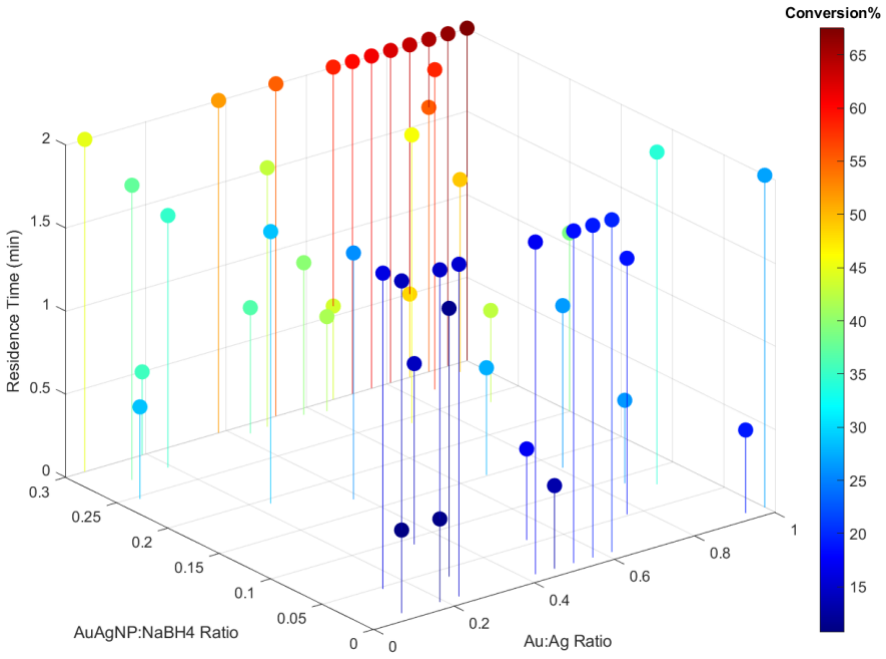

Figure 9: Response surface resulting from optimisation involving 21 linearly spaced levels of Au:Ag ratio in a nanoparticle catalyst.

N exponentially spaced levels (see equation 1):

5 levels of Au:Ag Ratio

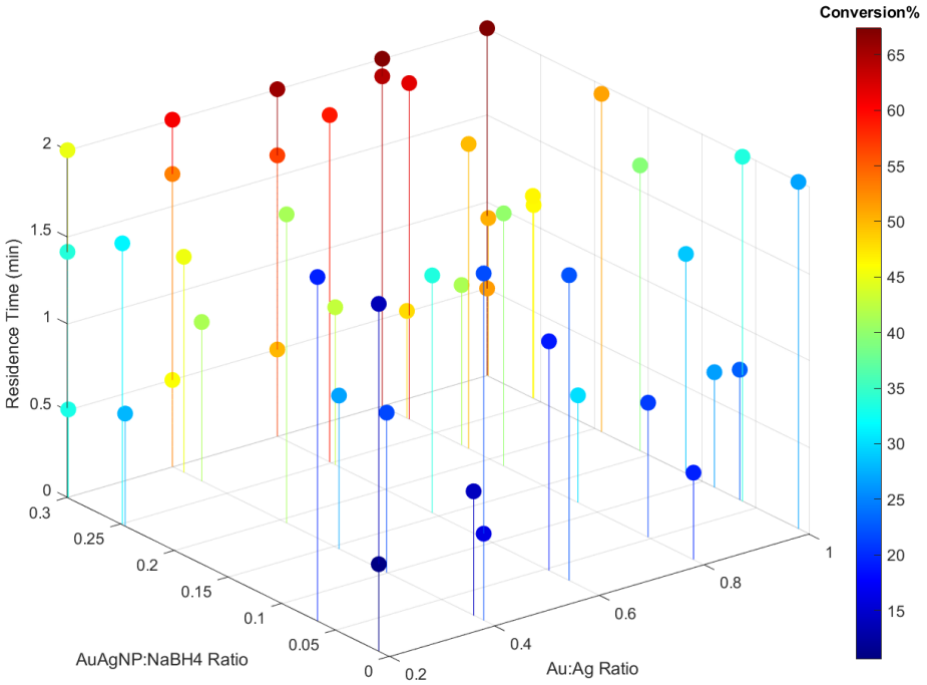

Figure 10: Response surface resulting from optimisation involving 5 exponentially spaced levels of Au:Ag ratio in a nanoparticle catalyst.

11 levels of Au:Ag Ratio

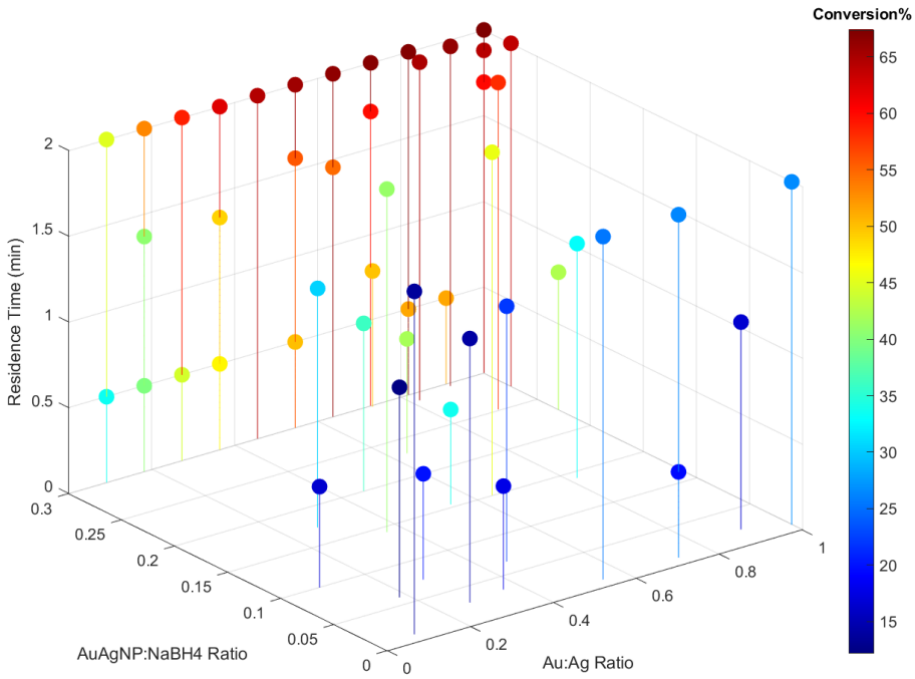

Figure 11: Response surface resulting from optimisation involving 11 exponentially spaced levels of Au:Ag ratio in a nanoparticle catalyst.

15 levels of Au:Ag Ratio

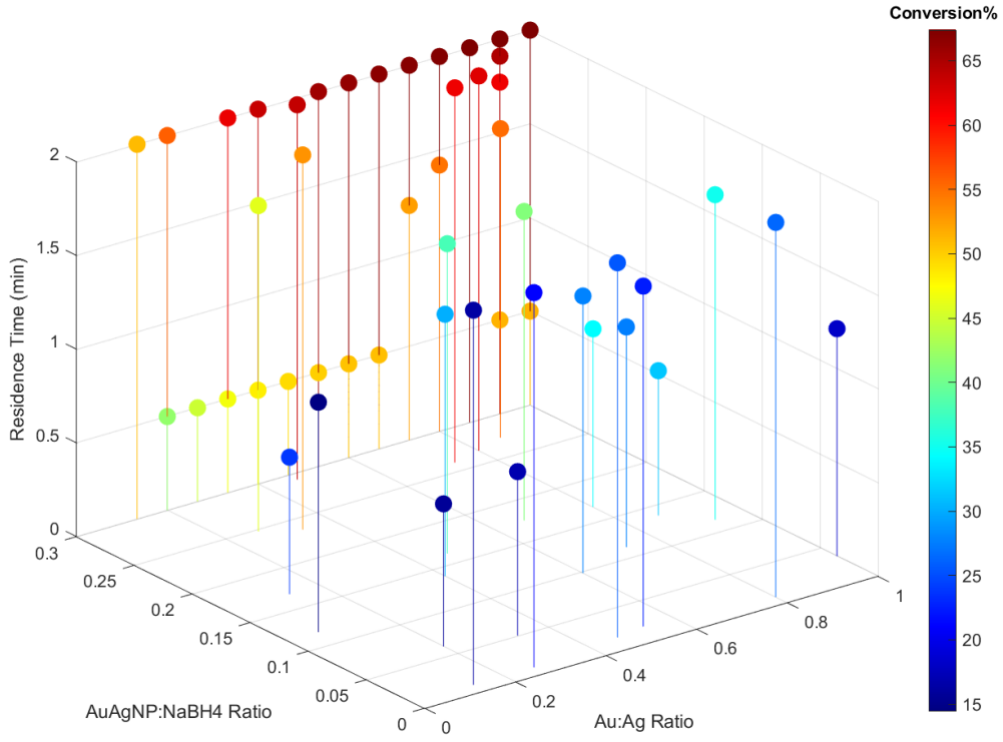

Figure 12: Response surface resulting from optimisation involving 15 exponentially spaced levels of Au:Ag ratio in a nanoparticle catalyst.

19 levels of Au:Ag Ratio

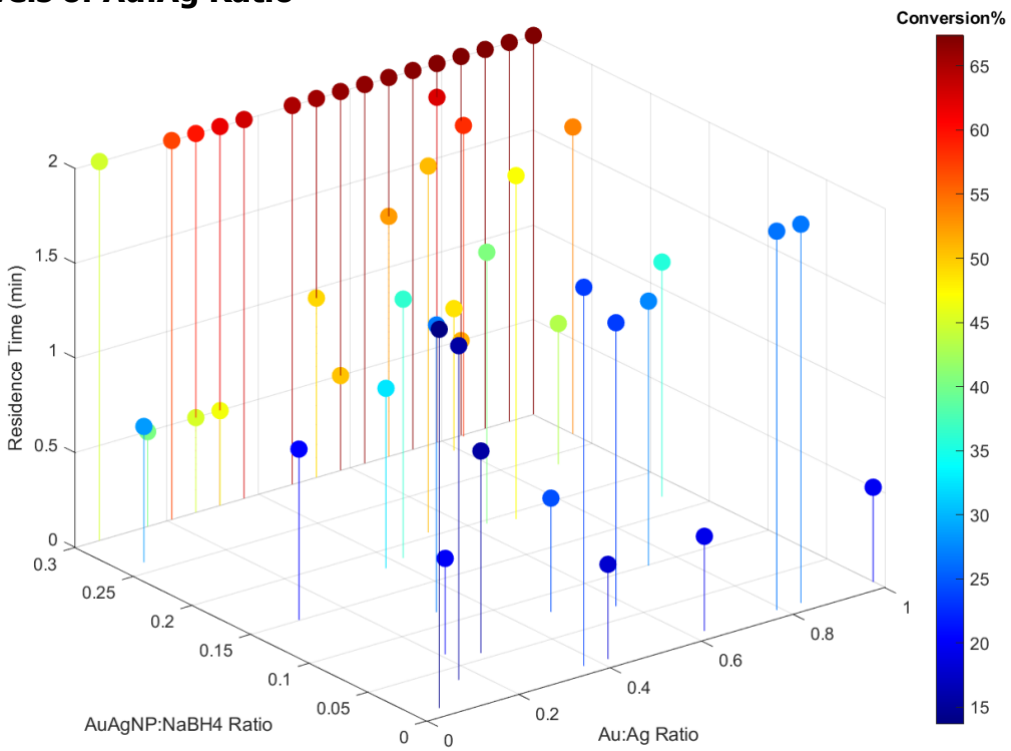

Figure 13: Response surface resulting from optimisation involving 19 exponentially spaced levels of Au:Ag ratio in a nanoparticle catalyst.

Table 2: Maximum conversion % per experiment

Experiment	Max. conversion (%)
B. Hall	72.9
Linearly Spaced (levels)	
5	67.6
11	67.6
15	67.6
21	67.6
Exponentially spaced (levels)	
5	67.4
11	67.4
15	67.4
19	67.4

Cumulative Conversion Percentage:

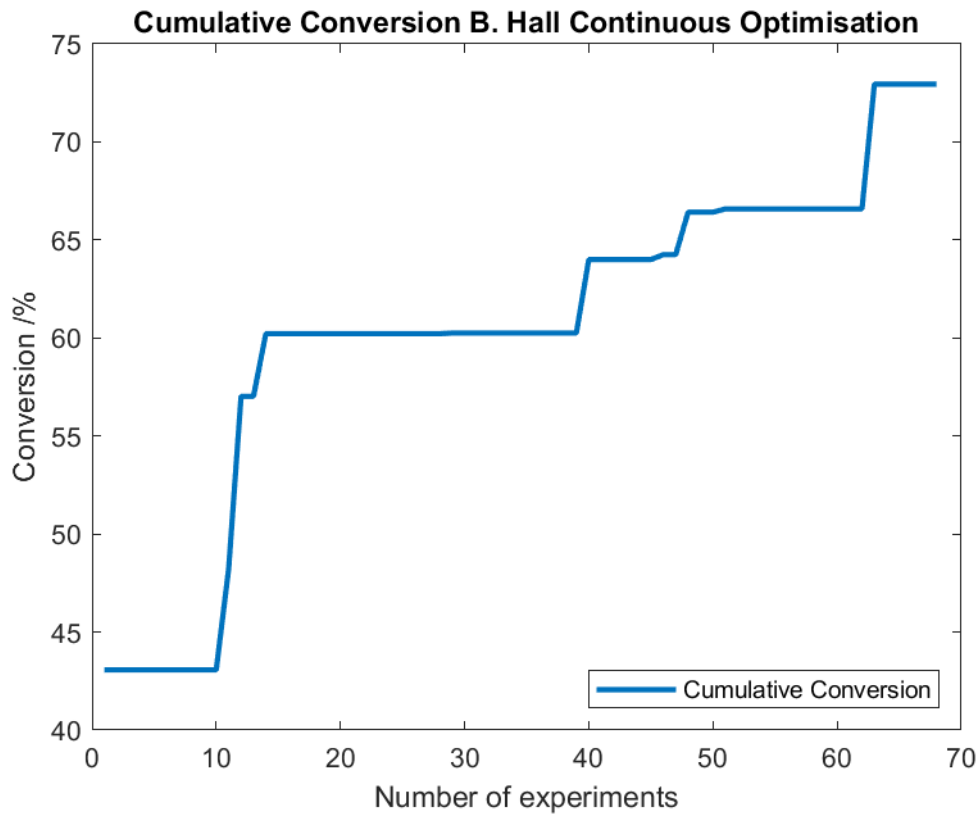

Figure 14: Cumulative conversion achieved by B. Hall in real-life continuous optimisation with the cont. flow reactor [6]

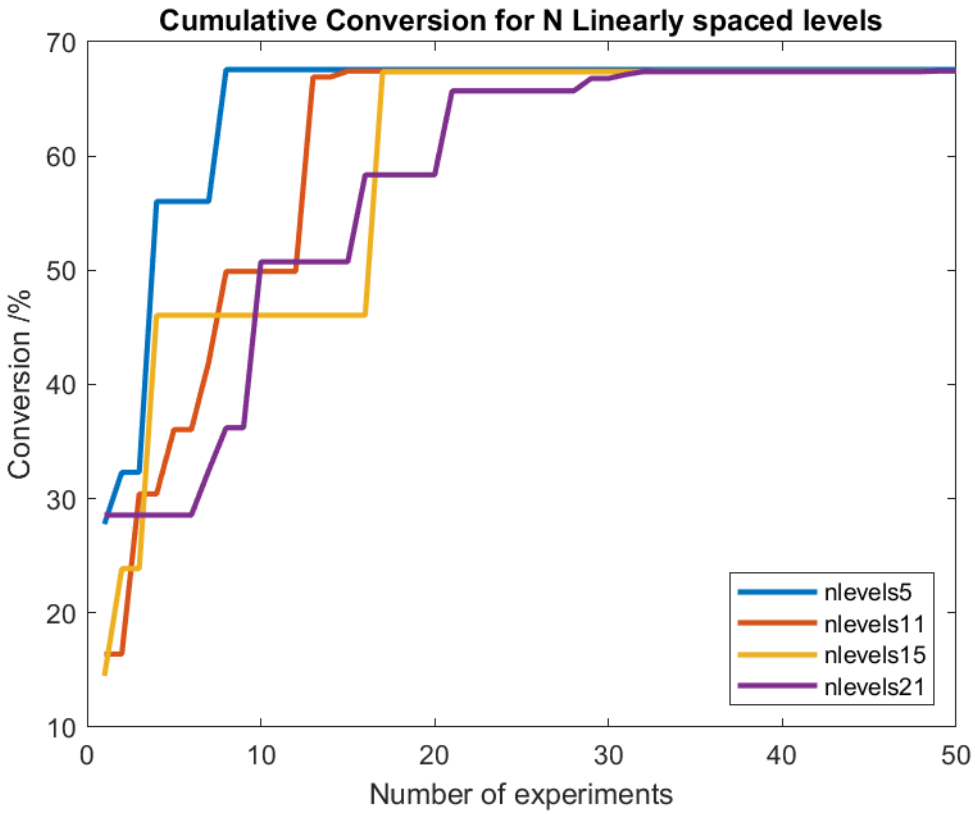

Figure 15: Cumulative conversion comparison (linearly spaced levels)

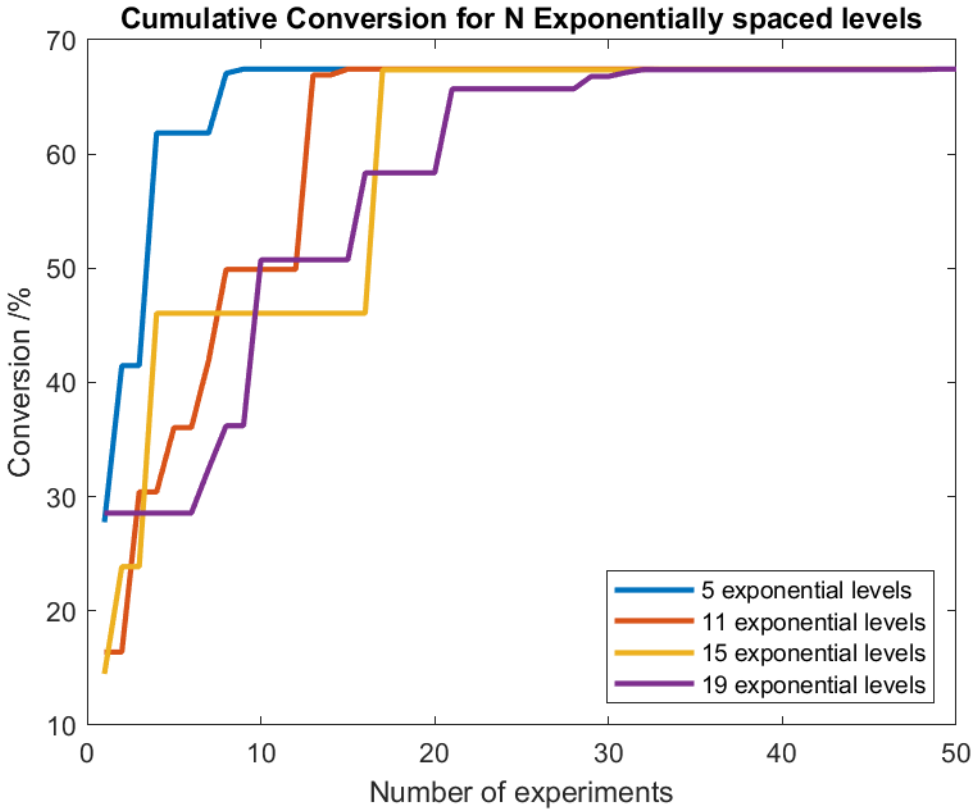

Figure 16: Cumulative conversion comparison (exponentially spaced levels)

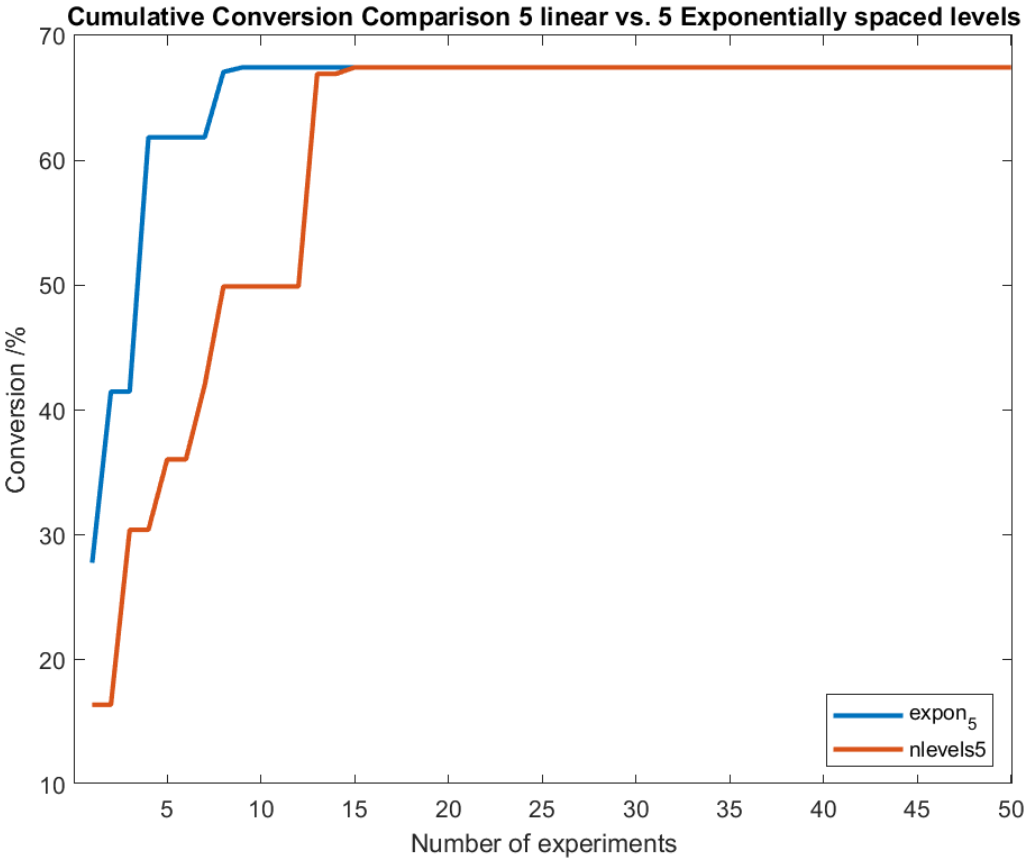

Figure 15: Cumulative conversion comparison (5 linear vs. 5 exponentially spaced)

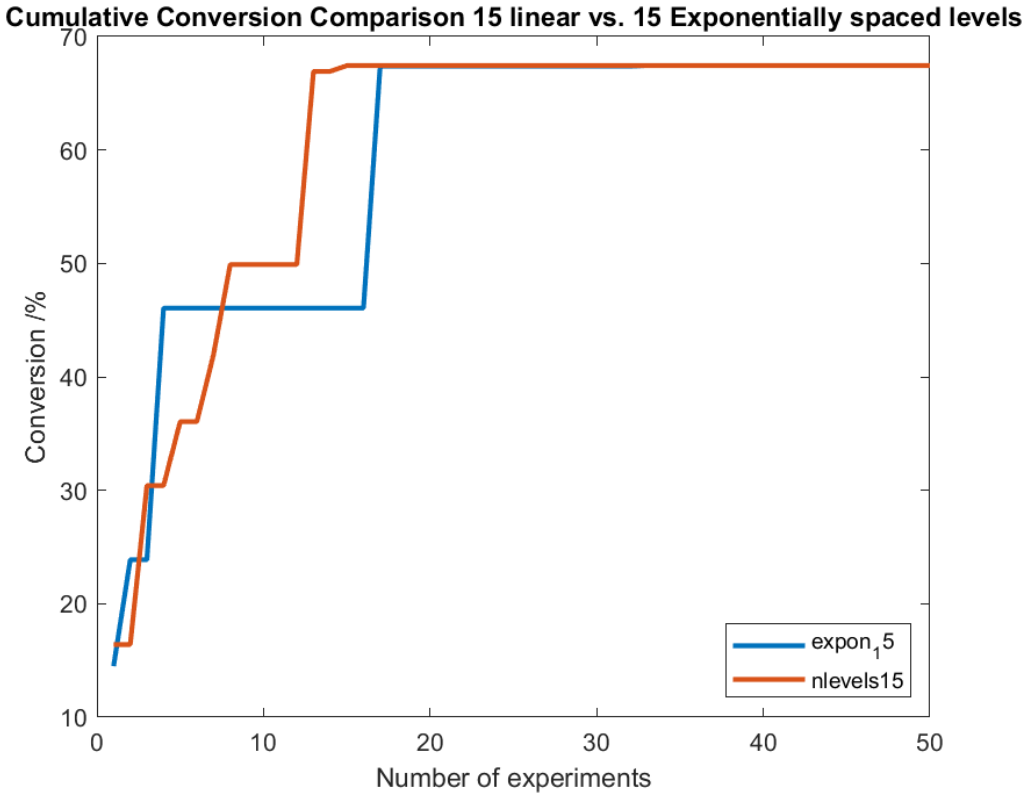

Figure 1816: Cumulative conversion comparison (15 linear vs. 15 exponential)

Exploration factor:

Table 3: Exploration factor calculated using the approach described by L. Power, of dividing the total distances between every point produced in each optimisation by the volume of experimental space (link to code in output).

Experiment	Exploration Factor
B. Hall	2700
5 linear levels	3323
11 linear levels	3137
15 linear levels	3633
19 linear levels	3588
5 exponential levels	3387
11 exponential levels	3522
15 exponential levels	3406
19 exponential levels	3447

This project returned mixed results. The first issue was that the discrete optimisation algorithm simulated experiments never achieved as high a conversion percentage as the continuous optimisation received by B. Hall. This provides evidence that the discrete algorithm was unable to access the necessary resolution in experimental space to find the true global optimum achieved in the continuous optimisation by B. Hall. Surprisingly, however, while the true optimum was never found, the optimum reached (~5% lower) was reached after significantly fewer experiments, shown by the cumulative conversion.

B. Hall took 65 experiments to reach the global optimum, in comparison to between 5 and 30 experiments with linearly spaced levels, and between 8 and 32 experiments for exponentially spaced levels.

While increasing the number of discrete levels in this system did not change the ultimate optimum produced, it did decrease the efficiency of the algorithm proportionally with the increased number of levels, such that the optimum took more experiments to achieve. This could be attributed to inaccurate tuning of the acquisition function leading to conservative exploration among the allowed discrete levels.

The exploration factor for the discrete optimisations were all slightly higher than the continuous by B. Hall, but this data is difficult to interpret considering the true optimum was missed. Further experimentation with this exploration factor would be necessary to determine real life effect of small differences in the factor between optimisations of differing types.

7. Conclusions & Future Work

Conclusions:

In a real-life system, these results would point towards better efficiency in terms of time and consumption of resources using the discrete approach, at the expense of an inaccurate global optimum reached. However, this can be attributed to the fact the system was not allowed to access the experimental space of the global optimum reach by B. Hall. This could be remedied by specifically setting the system to be able to access this space but would defeat the aims of the experiment of reducing the prior system knowledge requirement, which highlights the issues with the discrete approach.

As expected, the exponentially spaced levels achieved the optimum with fewer experiments than linear spaced, providing evidence that exponentially spaced levels in discrete optimisation are more efficient at producing the optimum in a "more equals better" system. However, this also requires prior knowledge, either from experiment or modelling, that it is indeed that kind of system. Using exponentially spaced levels in this system did not change the optimum reached, which was unexpected since doing so was supposed to circumvent the need for increased resolution specifically in this type of system; this should be investigated further in the future.

The exploration factor used in this project provides an insight into the amount of experimental space explored in a system, and a criticism is that it seems to only be comparable between similar systems due to the range of data, as well as the number of experimental data points. All the discrete optimisations showed greater exploration across the whole data set; despite this, the true optimum by B. Hall was unable to be achieved. This provides evidence that increased exploration alone is unable to make up for the reduced resolution in experimental space that discrete optimisation is restricted to.

This project does not provide definitive evidence that continuous variable optimisation is better than discrete. Discrete optimisation showed efficiency benefits over continuous, with the sacrifice of accuracy.
Due to the intense computational workload necessary for Bayesian optimisation, the higher numbers of discrete levels investigated beyond the results presented in this work were omitted due to concerns about validity. This shows that the discrete optimisation algorithm could be tuned to work more efficiently with higher numbers of discrete levels, perhaps utilising a more powerful computer than a home or office laptop.

Future work:

The results from this project point towards the need for refinement in the continuous optimisation system used in the continuous flow reactor by B. Hall. The approach used in this project could also be extended by zooming in on the higher end of the Au:Ag ratio, which due to technical issues with the optimiser, was unable to be achieved.
Results from this work would also point towards efficiency advantages in terms of time and resources of using discrete variable optimisation over continuous, but limitations in finding the 'true optima' of the system, at least in reaction systems like the ones used in this project.
Further work must also be done to test the validity of the exploration factor, and its ability to compare optimisations of different types.

8. Outputs, Data & Software Links

(Link to GitHub for exploration factor code:
https://github.com/louisgreenhalgh/matlabExplorationFactor.git)

(Link to poster presented at the Ai3SD Summer symposium 2021:
https://www.ai3sd.org/wp-content/uploads/sites/374/2021/08/11_Louis_Greenhalgh.pdf)

9. References

[1] R. Leardi, Anal. Chim. Acta, 2009, 652, 161–172.
[2] R. Lee, Chem. Ing. Tech., 2019, 91, 191–200.

[3] Keßler, T., Logist, F. and Mangold, M. (2016). Use of predictor corrector methods for multi-objective optimization of dynamic systems. Computer Aided Chemical Engineering, pp.313–318.

[4] Ganesan, Timothy; Elamvazuthi, Irraivan; Vasant, Pandian; Shaari, Ku Zilati Ku (2015-03-23). Nguyen, Ngoc Thanh; Trawiński, Bogdan; Kosala, Raymond (eds.). Multiobjective Optimization of Bioactive Compound Extraction Process via Evolutionary Strategies. Lecture Notes in Computer Science. Springer International Publishing. pp. 13–21. doi:10.1007/978-3-319-15705-4_2. ISBN 978-3-319-15704-7.

[5] Manson, J.A., Chamberlain, T.W. & Bourne, R.A. MVMOO: Mixed variable multi-objective optimisation. J Glob Optim 80, 865–886 (2021). https://doi.org/10.1007/s10898-021-01052-9

[6] B. Hall, Unpublished thesis: Self-optimising reactors for nanoparticle catalysis, University of Leeds (expected submission 2022)

[7] Bastus, N. G. et al. Kinetically Controlled Seeded Growth Synthesis of Citrate-Stabilized Gold Nanoparticles of up to 200 nm: Size Focusing versus Ostwald Ripening. Langmuir 27, 11098– 11105 (2011)