

Autonomous Experimentation: Active Learning for Enzyme Response Characterisation

Chris Lovell

Gareth Jones

Steve R. Gunn

Klaus-Peter Zauner

School of Electronics and Computer Science

University of Southampton, UK

CJL07R@ECS.SOTON.AC.UK

GJ07R@ECS.SOTON.AC.UK

SRG@ECS.SOTON.AC.UK

KPZ@ECS.SOTON.AC.UK

Editor: I. Guyon, G. Cawley, G. Dror, V. Lemaire, and A. Statnikov

Abstract

Characterising response behaviours of biological systems is impaired by limited resources that restrict the exploration of high dimensional parameter spaces. Additionally, experimental errors that provide observations not representative of the true underlying behaviour, mean that observations obtained from these experiments cannot be regarded as always valid. To combat the problem of erroneous observations in situations where there are limited observations available to learn from, we consider the use of multiple hypotheses, where potentially erroneous observations are considered as being erroneous and valid in parallel by competing hypotheses. Here we describe work towards an autonomous experimentation machine that combines active learning techniques with computer controlled experimentation platforms to perform physical experiments. Whilst the target for our approach is the characterisation of the behaviours of networks of enzymes for novel computing mechanisms, the algorithms we are working towards remain independent of the application domain.

Keywords: automatic hypothesis generation, closed-loop experimentation

1. Introduction

Nature exhibits biological systems that provide excellent computational mechanisms. Fundamental to that is the interactions of proteins, which can provide non-linear computational abilities (Zauner and Conrad, 2001). Understanding the behaviours exhibited by these interactions can only be achieved through physical experimentation. However, the scale and complexities of these domains mean that the number of experiments that can be performed is always heavily restricted in comparison to the size of the space being searched. Realistically, an experimenter may afford only a handful of experiments per parameter dimension. However, physical experimentation by nature implies that those experiments will produce observations with questionable accuracy. The variability of biological experimentation in particular, means some observations will be unrepresentative of the true underlying behaviour. As such, biological response characterisation exhibits the problems addressed by active learning, namely that learning must occur with the minimal number of observations as performing experiments is expensive (Cohn et al., 1996). To minimise experimentation costs whilst maximising information gain, an autonomous experimentation machine is in




Figure 1: Flow of experimentation between an artificial experimenter and an automated experimentation platform. A prototype of the lab-on-chip platform in development is shown.

development for biological response characterisation that combines active learning to reduce the number of experiments required to learn, with a resource efficient lab-on-chip automated experiment platform that minimises the volumes of reactants required per experiment. Autonomous experimentation is a closed-loop iterative process, where a computational system proposes hypotheses and actively chooses the next experiment, which is performed by an automated experimentation platform, with the results being fed back to the computational system, as illustrated in Figure 1. Here we consider the machine learning component, which is able to learn from a small number of actively chosen observations, where the observations are noisy and potentially erroneous. The lab-on-chip platform is currently in development (Jones et al., 2010).

The development of closed-loop autonomous experimentation machines is still in its infancy. Examples have existed within areas such as electro-chemistry (Żytkow et al., 1990), enzyme characterisation (Matsumaru et al., 2002) and identifying the functions of genes (King et al., 2004). Whilst another approach developed algorithms capable to guide an experimenter to rediscover the urea cycle (Kulkarni and Simon, 1990). Generally such systems, also described as computational scientific discovery systems in the literature, have applied more ad-hoc approaches to experiment design, with the exception of King et al. (2004) that considered a more mathematical active learning approach. However, both autonomous experimentation and many active learning techniques fail to address the problem of learning from only very small sets of experimental observations and that those observations may be unrepresentative of the actual behaviours that should be observed.

Presented here is a technique for producing likely response models from limited, noisy and potentially erroneous observations. To handle the uncertainty presented within this problem, a multiple hypotheses approach is utilised, where differing views about the validity of the observations are considered in parallel. In particular, in instances where an observation does not agree with a hypothesis, new hypotheses are created that consider the observation as both valid and erroneous in parallel, until further experimental evidence is obtained to determine whether it was the observation or the hypothesis that was invalid. Additionally a surprise based active learning technique is presented that guides experimentation to quickly identify the features of the behaviour under investigation, whilst highlighting erroneous observations.

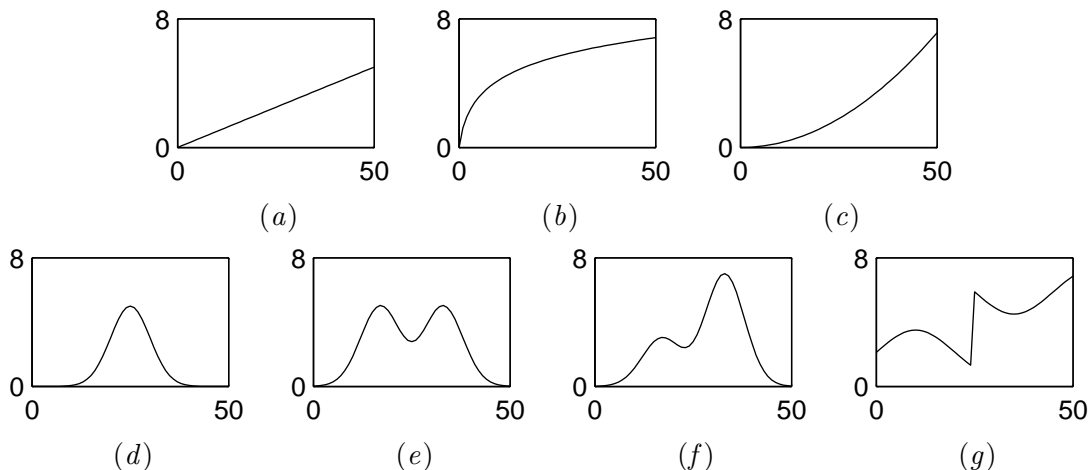


Figure 2: Underlying behaviours motivated from possible enzyme experiment responses.

2. Problem Formulation

The biological domains of interest currently do not have significant documented behaviours that can be used to validate the techniques proposed. Therefore to evaluate the approaches presented, we consider a generalised problem that closely matches the target problem domain. First we assume that the true underlying behaviour exhibited by the biological system under investigation, can be modelled by some function $f(x)$. The goal for the system is to build a function $g(x)$, which matches the response of $f(x)$. However, the responses from queries to $f(x)$ can be distorted by experiment measurement and reading errors, causing noise to be applied to both the responses (through ϵ) and to the requested experiment parameters (through δ). Additionally, the lack of control of the biological materials, also present distortions to the responses of $f(x)$. In enzyme experimentation, the reactants can undergo undetectable physical or chemical change, which leads to experiments with those reactants yielding erroneous observations, unrepresentative of the true underlying behaviour. We model such instances as shock noise (through ϕ), which applies a large offset to the response value. Whilst ϵ and δ can occur on every experiment, ϕ will only be non-zero for a small proportion of experiments. We do not consider the case where ϕ occurs for a large number of experiments, as in this instance, the results from such experimentation would be disregarded from consideration anyway. We therefore represent a response characterisation experiment as:

$$y = f(x + \delta) + \epsilon + \phi \quad (1)$$

where parameter x and response y can be replaced with vectors for higher dimensionality.

2.1. Underlying Behaviours

Whilst models of existing behaviours do not currently exist for the domain of interest, we can define some properties of those behaviours that may be expected or would be potentially useful for engineering with these biomolecules. In Figure 2, a range of underlying behaviours, f_a, \dots, f_g , are presented. These behaviours test, in figure order (a-g): (a) linear

response, (b) non-linear response, (c) power law, (d) single peak, (e) two peaks, (f) two peaks where one peak is dominant over the other, (g) discontinuity between two distinct behaviours. Behaviours (a-c) are motivated from expectations that behaviours are often described in terms of linear systems or power laws, where (b) is similar to Michaelis-Menton kinetics (Nelson and Cox, 2008) and (c) is similar to responses where there is a presence of cooperativity between substrates and enzymes (Tipton, 2002). Behaviours (d-g) are motivated from the belief that expected behaviours in the domain being investigated may be nonmonotonic and could also include a phase change between distinct behaviours (Zauner and Conrad, 2001). We next discuss the implementation issues of the computational side of autonomous experimentation.

3. Hypothesis Management

A key problem for a hypothesis manager, is how to handle uncertainty in the form of erroneous observations. By accepting all observations as valid, errors can mislead the development of hypotheses. Determining the validity of observations is impeded by the limited resources, which prevent repeat experiments. In this situation, maintaining a single hypothesis appears inefficient in obtaining an accurate representation of the underlying behaviour. Alternatively, we can consider using multiple hypotheses that maintain different views of the validity of the observations in parallel. Whilst many multiple hypotheses based approaches produce hypotheses using random subsets of the data (Freund et al., 1997; Abe and Mamitsuka, 1998), we believe a more structured approach can be applied to deal with the uncertainty about the validity of the observations. That is, where an observation appears erroneous, separate hypotheses can be used in parallel that consider the observation as erroneous or valid, with further experimentation providing the evidence to differentiate between the hypotheses.

To illustrate this, consider the situation presented in Figure 3, where observations are labelled alphabetically in the order obtained. After the first two observations are obtained, hypothesis h_1 appears as a reasonable hypothesis. On obtaining observation C however, a potential flaw in this hypothesis is found, suggesting that the hypothesis is erroneous, or with the expectation of erroneous observations, the observation itself could be erroneous. Continuing with the acquisition of observation D, the validity of observation C is now more likely, however observation B is now of questionable validity.

To achieve different views of observations with questionable validity, observations can be weighted differently in the regression calculation. In Zembowicz and Żytkow (1991) and Christensen et al. (2003), where the accuracy of observations is known, deliberate weighting of observations has been applied to obtain better predictions of the underlying behaviours. But in the present problem, obtaining accuracy information is restricted by resources. Multiple hypotheses allows different views about the validity of the observations to be considered in parallel, allowing any decisions about observation validity to be postponed until sufficient evidence is available. In the following section we describe this process in more detail.




Figure 3: Validity of observations affecting hypothesis proposal. Hypotheses (lines) are formed after observations (crosses) are obtained. In (a), h_1 formed after A and B are obtained questions the validity of C. In (b), D looks to confirm the validity of C however causes h_4 and h_5 to differ in opinion about the validity of B.

3.1. Implementation

In practice a hypothesis is represented here by a smoothing spline. A smoothing spline is a piecewise cubic spline regression technique that can be placed within a Bayesian framework (Wahba, 1990):

$$S_{w,\lambda}(f) = \sum_{i=1}^n w_i (y_i - f(x_i))^2 + \lambda \int_a^b (f''(x))^2 dx \quad (2)$$

where experiment parameter and observation pairs x_i and y_i are used to train a regression fit of the data. The parameter w_i is a weighting applied to each x_i, y_i pair, and the hyperparameter λ controls the amount of regularisation, with b and a being the maximum and minimum of the x_i values respectively.

The w and λ parameters are chosen by the hypothesis manager for each hypothesis, the method for this follows, such that a hypothesis, h , is the minimiser of the smoothing spline regression function for a particular w and λ :

$$h = \min S_{w,\lambda}(f) \quad (3)$$

The process of the hypothesis manager is as follows. After an experiment has been performed, a set of new hypotheses are proposed. New hypotheses are created from random subsets of the available observations, along with a randomly selected smoothing parameter, so as to allow for different initial views of the parameter space. All new hypotheses are added to the set of working hypotheses. The smoothing parameter is chosen from a set of possible parameters ($\lambda \in \{10, 50, 100, 150, 500, 1000\}$) that allow for a range of different fits of the data, corresponding to different initial views of the behaviour being investigated.

Next the hypothesis manager reviews the validity of the observations that have been obtained. To do this, the hypothesis compares all observations against all of the working hypotheses. Through the smoothing spline, each hypothesis is able to provide an indicative error bar for the prediction of the outcome of a particular experiment parameter (Wahba,

1990). This error bar is used to determine whether or not an observation agrees with a hypothesis, where if the observation falls outside of the error bar value, the observation is said to be in disagreement with the hypothesis. When such a disagreement occurs between a hypothesis and an observation, all of the parameters for that hypothesis are taken, and used to build two new refined hypotheses. These refined hypotheses differ from the original hypothesis, through altering the weighting parameter applied to the observation of questionable validity. One hypothesis will set the weighting applied to the observation to be 0 and the other to 100, so as to create a hypothesis that considers the observation to be erroneous and another that considers the observation to be true, where the high weight will force the outcome of the regression to pass closer to the observation. These two new hypotheses along with the original hypothesis, are kept in the working set of hypotheses.

After this process of refinement, all hypotheses in consideration are evaluated against the available observations using the following function:

$$C(h) = \frac{1}{N} \sum_{n=1}^N \exp \left(\frac{-\left(\hat{h}(x_n) - y_n\right)^2}{2\sigma^2} \right) \quad (4)$$

where $\hat{h}(x_n)$ is the hypotheses prediction for experiment parameter x_n , with y_n being the experimental observation for parameter x_n , σ is chosen a priori (currently 1.96), and N is the number of observations. Finally, for computational efficiency, the number of working hypotheses considered in parallel can be reduced. Removing the hypotheses that perform poorly in the evaluation stage, ensures that whilst the number of hypotheses considered in parallel remains large, it does not become computationally infeasible to inspect in the experiment selection stage. In the trials presented in Section 5, 200 new random hypotheses are created in each iteration, and the best 20% of all hypotheses under consideration are maintained into the next round of experimentation. From initial trials it appears that so long as the number of new hypotheses created is large, the number of hypotheses retained after each experiment can be altered as required for performance.

In review, the hypothesis manager maintains an expanding ensemble of working hypotheses throughout the experimentation conducted, where the above process of hypothesis proposal is conducted after every experiment performed. Hypotheses have parameters for the observation weightings, w_i , and smoothing parameter, λ . When creating new hypotheses, the hypothesis manager chooses initial random parameters, through selecting a random subset of the available observations to train from, giving those observations all initial weights of 1, with the rest 0, along with a randomly selected smoothing parameter. The parameter learning for the hypotheses comes through the refinement of the existing hypotheses, where the weight parameters of an existing hypothesis are changed in the new refined hypothesis to either 0 or 100, depending on whether the observation is believed to be erroneous, or valid but indicating a feature of the behaviour not characterised in the original hypothesis. The original hypothesis and the subsequent refinements are then all maintained in the working set of hypotheses, so as to test the new parameter settings in parallel. Only when sufficient experimental evidence is available that contradicts a particular hypothesis, is that hypothesis along with its set of parameters removed from consideration, whilst the more suitable hypotheses remain. Next we discuss how the set of hypotheses in consideration can be used to provide information for determining the experiments to perform.

4. Active Learning Experiment Management

The role of the experiment manager is to employ active learning techniques, to determine the next experiments to perform. The experiment manager uses the information available to it, namely the observations obtained and the hypotheses under consideration. With the hypothesis manager providing a set of competing hypotheses, the experiment manager adopts a query-by-committee style approach for determining the experiments to perform. In query-by-committee, labels or observations as referred to here, are chosen where the committee members most disagree (Seung et al., 1992). In other words, the experiment manager should select the experiments that are most likely to differentiate between and in turn disprove hypotheses under consideration, which agrees the experimental design methods suggested in philosophy of science literature (Chamberlin, 1890).

Experimental design T-optimal approaches exist for separating sets of hypotheses, however they can perform poorly if there is experimental noise (Atkinson and Fedorov, 1975). Alternatively, ensembles of hypotheses have been differentiated between by placing experiments where the variance of the predictions of the hypotheses is greatest (Burbidge et al., 2007). However, selecting experiments where the variance of the hypotheses predictions is greatest, can be misled by outlying hypothesis predictions, as shown in Figure 4.

Therefore we require an alternative active learning technique that will separate hypotheses efficiently. To achieve this we consider a strategy that chooses experiments where there is the maximal disagreement between any two of the hypotheses under consideration:

$$D = \arg \max_x \sum_{i=1}^N \sum_{j=1}^N \int (P_{h_i}(y|x) - P_{h_j}(y|x))^2 dy \quad (5)$$

By replacing the y integral with the prediction of a hypothesis, $\hat{h}(x)$, the following equation will separate a set of hypotheses based on their predictions for different x :

$$D'(x) = \sum_{i=1}^n \sum_{j=1}^n 1 - \exp \left(\frac{-\left(\hat{h}_i(x) - \hat{h}_j(x)\right)^2}{2\sigma_i^2} \right) \quad (6)$$

where \hat{h}_i is the prediction of hypothesis i and σ_i^2 comes from the error bar of h_i for x . This discrepancy approach is more robust than a variance method, as shown in the example set of hypotheses shown in Figure 4, where a variance method would place an experiment where the prediction of the majority of the hypotheses is the same. Further to this, this discrepancy can adapt with the previous observations available, so as to differentiate between only the well performing and currently agreeing hypotheses:

$$D(x) = \sum_{i=1}^n \sum_{j=1}^n C(h_i)C(h_j)A(h_i, h_j) \left(1 - \exp \left(\frac{-\left(\hat{h}_i(x) - \hat{h}_j(x)\right)^2}{2\sigma_i^2} \right) \right) \quad (7)$$

where $A(h_i, h_j)$ is the agreement between the hypotheses for the previous experiments:

$$A(h_i, h_j) = \frac{1}{N} \sum_{n=1}^N \exp \left(\frac{-\left(\hat{h}_i(x_n) - \hat{h}_j(x_n)\right)^2}{2\sigma_i^2} \right) \quad (8)$$




Figure 4: Location of experiments selected to maximise discrepancy between hypotheses. Solid bold vertical line is the experiment parameter the variance approach chooses. Dashed bold vertical line is the experiment parameter the maximum discrepancy approach in Equation (6) chooses. The curves show the predictions of the hypotheses across the parameter space.

which could alternatively be calculated as a product of the agreement.

This discrepancy equation exploits the hypotheses to provide experiments that can effectively discriminate a set of hypotheses. However, it does not explore the experiment parameter space, which is needed to allow the hypothesis manager to build representative hypotheses in the first place. This can in part be addressed by performing an initial number of exploratory experiments, which also allows for the first hypotheses to be proposed. However, additional consideration must be given to handle this exploration-exploitation trade-off (Auer, 2002). In the following sections, two new experiment selection techniques are presented that consider this trade-off.

4.1. Exploring Peaks in the Discrepancy Equation

By placing experiments where $D(x)$ is maximal, experiments may end up being placed within the same localised area of the experiment parameter space, repeatedly investigating one particular discrepancy, without any exploration. However, if we consider $D(x)$ over all possible experiment parameters, there will likely be local maxima, or peaks, in different areas of the parameter space. These local maxima show different features of the behaviour where the hypotheses disagree elsewhere in the parameter space. Therefore, instead of selecting the maximum of $D(x)$, the maxima can be used to select a set of experiments to perform across the parameter space, which investigate different reasons for hypothesis disagreement, whilst simultaneously allowing some additional exploration.

The process for this experiment selection technique is as follows. Starting with the initial observations and hypotheses, a set of experiments to perform are chosen as those at the peaks of $D(x)$, where experiments are not repeated. Those experiments are then chosen in order of their $D(x)$ value, from largest to smallest, so that if resources are depleted, then the experiments that are likely to differentiate between the hypotheses the most, will have been performed. After each experiment is conducted, new hypotheses are created, but the next set of experiments to perform are only chosen once the current set of experiments have

been performed. This process continues until the maximum allowed number of experiments determined by the user have been performed.

4.2. Surprise Based Exploration-Exploitation Switching

Investigating surprising observations, defined as those observations that disagree with a well performing hypothesis, has been highlighted as a technique utilised by successful human experimenters and has also been considered in previous computational scientific discovery techniques (Kulkarni and Simon, 1990; Matsumaru et al., 2002). A surprising observation either highlights a failure in the hypothesis or an erroneous observation. If the observation is highlighting a failure of a hypothesis, especially an otherwise well performing hypothesis with a high prior confidence, then additional experiments should be performed to further investigate the behaviour where that observation was found, to allow the development of improved hypotheses. As such we consider the use of surprise to manage the exploration-exploitation trade-off, where obtaining surprising observations will lead to more exploitation experiments, and unsurprising observations lead to exploration experiments.

A Bayesian formulation for surprise has been considered previously in the literature, where a Kullback-Leibler divergence is used to identify surprising improvements to the models being formed (Itti and Baldi, 2009). However, the surprise in Itti and Baldi (2009) is scaled by higher posterior probabilities, but here we are more interested in those hypotheses with high prior confidences but lower posterior confidences as a result of the last experiment. Whilst looking for reductions in posterior probability may appear counter-intuitive, it is important to remember that successful refinement of those hypotheses will result in new hypotheses with higher confidences. Therefore, we interchange the prior and posterior terms to rework the Bayesian surprise function to be:

$$S = \sum_i C(h_i) \log \frac{C(h_i)}{C'(h_i)} \quad (9)$$

where $C(h)$ is the prior confidence of h before the experiment is performed, and $C'(h)$ is the posterior confidence of h after the experiment has been performed, calculated across all hypotheses under consideration using Equation (4), before any new hypotheses are added. Positive values of S states that the observation was surprising, as the overall confidence of the hypotheses have been reduced. Whilst a negative value states the observation was not surprising, as the overall confidence has increased. The result of S can therefore be used to control the switching between exploration and exploitation experiments, where a positive value will dictate that the next experiment will be exploitative, so as to allow investigation of the surprising observation. Whilst a negative value of S will lead to an exploration experiment next, to search for new surprising features of the behaviour.

The procedure for this experiment selection technique is as follows. The prior confidence of the current set of hypotheses before the experiment is performed, is compared with the posterior confidence of those same hypotheses after the experiment is performed, using the surprise function of Equation (9). If $S > 0$ then an exploitation experiment, the maximum of the discrepancy equation $D(x)$, will be performed on the next iteration. Otherwise an exploration experiment will be performed, which is defined as the experiment that has the maximum minimum distance to any previously performed experiment in the experiment

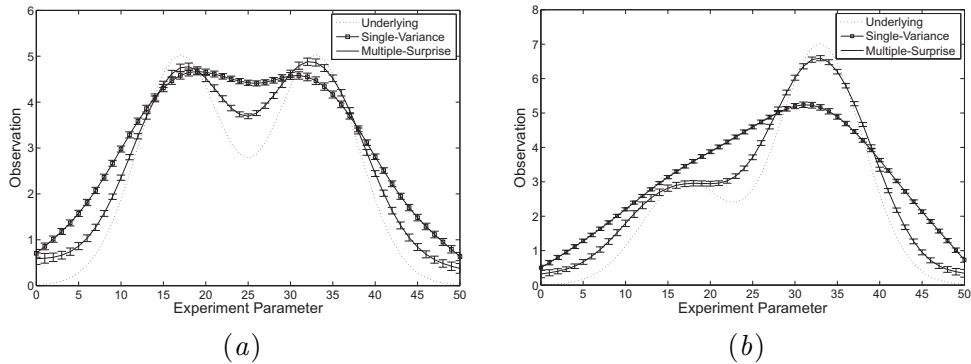


Figure 5: Comparison between the true underlying behaviour and the mean of the most confident hypotheses predictions for 100 trials, for the single hypothesis approach using prediction variance experiment selection, and the multiple hypotheses approach using the surprise technique for selecting exploration or exploitation experiments. Shown using behaviour f_e (a) and f_f (b).

parameter space. After S has been calculated, the hypothesis manager will go through the process of creating new hypotheses. This process of evaluating experiments using surprise to choose the next experiment type, is continued until the maximum number of experiments allowed has been performed.

5. Results and Discussion

Simulated experiments are conducted using the behaviours described in Section 2.1. All observations have additional Gaussian noise $\epsilon = N(0, 0.5^2)$. Parameter shift noise is kept here at $\delta = N(0, 0)$, for clarity of results presented here, as such noise in initial trials appears to have little impact in the performance of the approaches tested here. Experiments are bounded between 0 and 50, and are discretised evenly over the parameter space with 51 possible different experiments, first to make experiment selection more tractable, but also to reflect that physical experiment parameter spaces have finite precision controlled by the laboratory hardware available. Initially 5 exploration experiments are performed that are equidistant to one another in the parameter space, to allow for an initial set of hypotheses to be proposed. One of these initial experiments in each trial has random shock noise $\phi = N(3, 1)$ applied to it. The evaluation of the techniques occur over 15 actively selected experiments, where 3 of those experiments produce erroneous observations.

To contrast the multiple hypothesis approach, a single hypothesis approach is used. The single hypothesis is trained with all available observations, using cross-validation to determine the smoothing parameter. Experiments are chosen in the single hypothesis case through random selection, or where the error bar of the hypothesis is greatest. The multiple hypotheses method has experiments chosen through: random selection; choosing the maximum discrepancy value; choosing the peaks of the discrepancy function; and using the surprise method to switch between exploration and exploitation experiments.

AUTONOMOUS EXPERIMENTATION

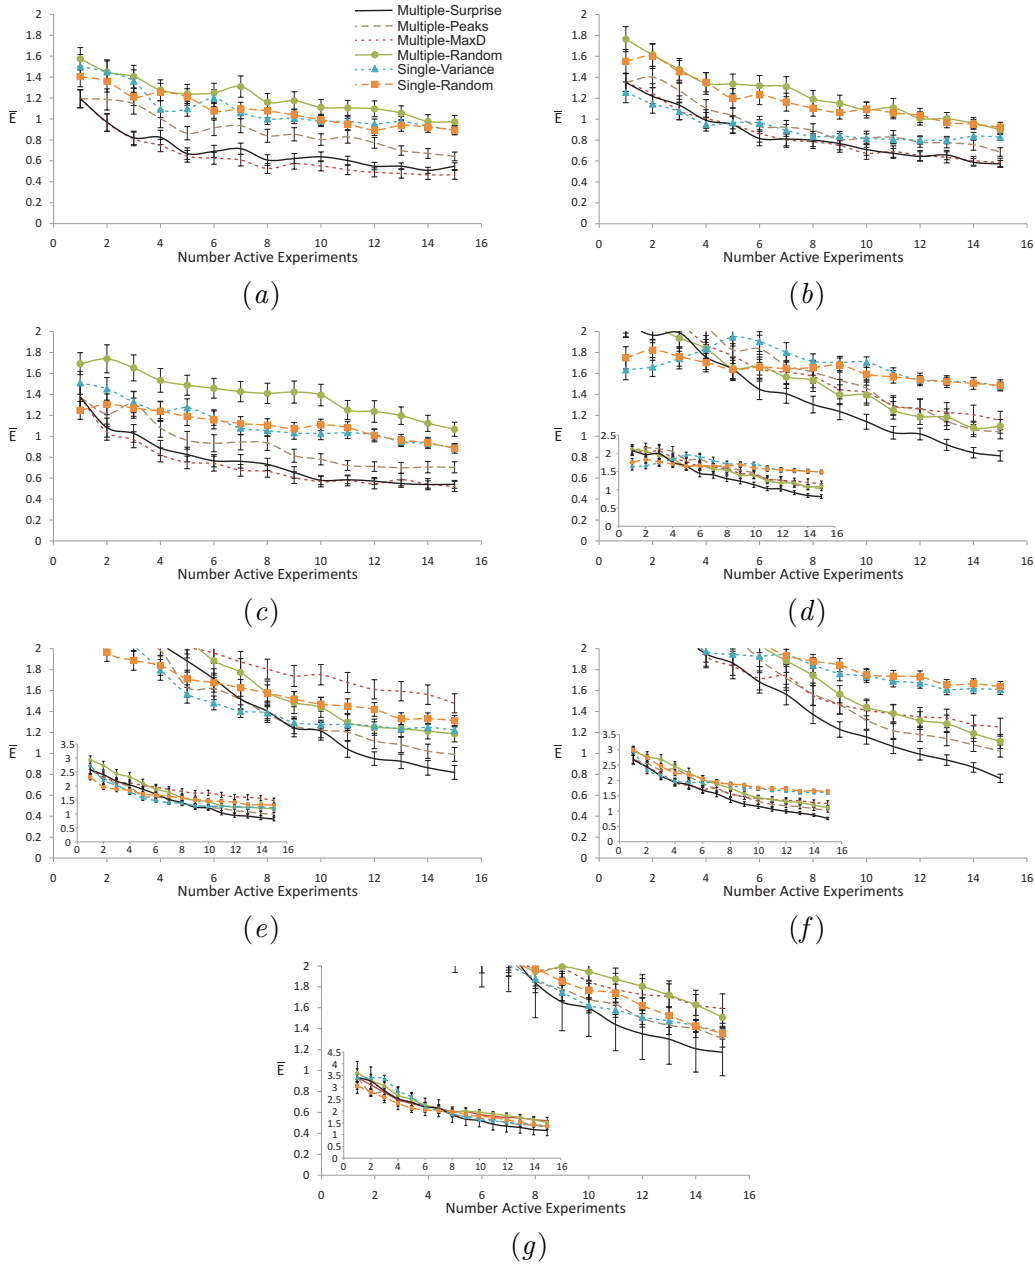


Figure 6: Performance of active learning and hypothesis management techniques. Shown is a comparison of error between the most confident hypothesis and the true underlying behaviour, over the number of actively chosen experiments, where 20% of the observations are erroneous, for 100 iterations. Shown in (a–g) are the corresponding results for the 7 behaviours shown in Figure 2.

To evaluate, the mean squared error between the most confident hypothesis of each trial is compared to the underlying behaviour being investigated:

$$E = \frac{1}{N} \sum_{n=1}^N \left(\hat{b}(x_n) - f(x_n) \right)^2 \quad (10)$$

where $\hat{b}(x_n)$ is the prediction of the most confident hypothesis in the trial, for parameter values chosen across the whole parameter space. The mean of these trials over 100 iterations for the 7 different underlying behaviours considered, is shown in Figure 6.

Throughout, the single hypothesis techniques perform poorly in comparison to the multiple hypothesis techniques. Poor performance is due to the single hypothesis generally averaging through all of the data, which can result in features of the behaviours being missed, especially in the more complex nonmonotonic behaviours (d–g), as shown in Figure 5. In the monotonic cases, the difference in performance between the single and multiple hypotheses techniques comes from the single hypothesis averaging through all observations, including the erroneous ones, which allows the erroneous observations to affect the predicted responses, making the hypothesis less accurate.

The multiple hypotheses techniques generally outperform the single hypothesis methods, however the extent of which is dependent on the active learning technique employed. The random strategy performs poorly in the monotonic behaviours (a–c), as experiments are not performed specifically to evaluate the accuracy of observations, which allows for the hypotheses to be misled by the erroneous observations. Whilst this is still an issue in the nonmonotonic behaviours (d–g), the random strategy will generally explore the parameter space more, so identifying the different features of the behaviour being investigated, leading it to have a lower error rate than the single hypothesis techniques, and occasionally similar to the other multiple hypotheses techniques. The maximum discrepancy technique (MaxD) performs well in the simpler monotonic behaviours, as most of the differences between hypotheses will be caused by erroneous observations, which the technique will investigate and be able to produce an accurate representation of the behaviour. In the monotonic behaviours however, the technique may miss some of their features, where its success in identifying the features is dependent on the initial exploratory experiments, as it will perform no exploration on its own and may become stuck investigating the same feature repeatedly. Using the peaks of the discrepancy equation provides more exploration of the parameter space than choosing just the maximum of the equation, allowing for lower error values in the non-monotonic behaviours. However, in the monotonic behaviours the strategy may spend more experiments investigating small differences between the hypotheses than investigating erroneous observations, meaning that the resultant hypotheses are not as accurate as using the maximum discrepancy for these behaviours. The surprise technique performs consistently well for all behaviours tested, by being able to evaluate the accuracy of the observations and suitability of the hypotheses through exploitation experiments, whilst performing a small number of additional exploratory experiments to further investigate the parameter space.

Over the 100 trials, the surprise technique used few exploration experiments per trial, with an average of 5 exploration experiments in the monotonic cases, normally in the latter stages of experimentation, and 4 exploration experiments in middle to latter stages for the nonmonotonic cases. As the hypotheses quickly produce a good representation of the underlying phenomena in the monotonic cases, additional exploratory experiments are performed as the observations obtained are not surprising to the hypotheses. If we allow the multiple peaks technique to have an additional 5 initial exploratory experiments but with 5 fewer exploratory experiments, we find that it has a similar performance to the surprise method with only the 5 initial exploratory experiments, except for a significant improvement in predicting f_g by the multiple peaks technique. However, this is due to the initial 10 exploratory

experiments covering all features of the behaviour. The surprise technique is therefore more preferable than the multiple peaks technique, as it has a lower initial exploratory experiment requirement, instead deciding for itself whether additional exploration is required. As such the technique could be adapted to terminate experimentation after performing several unsurprising experiments, reducing the resources used further.

6. Conclusion

Presented is work towards an application for active learning, called autonomous experimentation. Our target domain is automatic enzyme characterisation, where the number of available experiments will be limited to a handful per parameter dimension and that those observations may be erroneous and unrepresentative of the true underlying behaviours. Our belief is that the uncertainty that exists within this problem, is best dealt with through a multiple hypotheses approach. In such an approach, decisions about the validity of observations can be delayed until more experimental evidence is available, through competing hypotheses with different views about the validity of the observations. These multiple hypotheses can be used for effective response characterisation when coupled to an active learning technique, which will outperform a single hypothesis based approach. A technique has been presented that evaluates the surprise of the previous experiment to determine whether the system will next perform an experiment that will explore the parameter space to find new features of the behaviour not yet represented by the hypotheses, or perform an experiment to exploit information held in the hypotheses so as to discriminate between them. The weakness of the multiple hypotheses technique has been shown to be where it is coupled with a random experiment strategy, where erroneous observations can be accepted as true, without experiments testing their validity, leading to the most confident hypotheses being created on inaccurate data. Our next step is to connect the algorithms presented here, with the microfluidic experimentation platform in development (Jones et al., 2010), to demonstrate fully autonomous experimentation.

Acknowledgments

This work was supported in part by a Microsoft Research Faculty Fellowship to KPZ.

References

- N. Abe and H. Mamitsuka. Query learning strategies using boosting and bagging. In *ICML '98*, pages 1–9, San Francisco, CA, USA, 1998. Morgan Kaufmann.
- A. C. Atkinson and V. V. Fedorov. The design of experiments for discriminating between several models. *Biometrika*, 62(2):289–303, 1975.
- P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- R. Burbidge, J. J Rowland, and R. D King. Active learning for regression based on query by committee. In *IDEAL 2007*, pages 209–218. Springer-Verlag, 2007.

- T. C. Chamberlin. The method of multiple working hypotheses. *Science (old series)*, 15: 92–96, 1890. Reprinted in: *Science*, v. 148, p. 754–759, May 1965.
- S. W. Christensen, I. Sinclair, and P. A. S. Reed. Designing committees of models through deliberate weighting of data points. *JMLR*, 4:39–66, 2003.
- D. A. Cohn, Z. Ghahramani, and M. I. Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996.
- Y. Freund, H. S. Seung, E. Shamir, and N. Tishby. Selective sampling using the query by committee algorithm. *Machine Learning*, 28:133–168, 1997.
- L. Itti and P. Baldi. Bayesian surprise attracts human attention. *Vision Research*, 49: 1295–1306, 2009.
- G. Jones, C. J. Lovell, H. Morgan, and K.-P. Zauner. Characterising enzymes for information processing: Microfluidics for autonomous experimentation (abstract). In *9th International Conference on Unconventional Computation*, page 191, Tokyo, Japan, 2010.
- R. D. King, K. E. Whelan, F. M. Jones, P. G. K. Reiser, C. H. Bryant, S. H. Muggleton, D. B. Kell, and S. G. Oliver. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:247–252, 2004.
- D. Kulkarni and H. A. Simon. Experimentation in machine discovery. In J. Shrager and P. Langley, editors, *Computational Models of Scientific Discovery and Theory Formation*, pages 255–273. Morgan Kaufmann Publishers, San Mateo, CA, 1990.
- N. Matsumaru, S. Colombano, and K.-P. Zauner. Scouting enzyme behavior. In D. B. Fogel et al., editor, *WCCI'02 – CEC*, pages 19–24, Honolulu, Hawaii, 2002. IEEE.
- D. L. Nelson and M. M. Cox. *Lehninger Principles of Biochemistry*. W. H. Freeman and Company, New York, USA, 5th edition, 2008.
- H. S. Seung, M. Opper, and H. Sompolinsky. Query by committee. In *Proceedings of the ACM Workshop on Computational Learning Theory*, pages 287–294, 1992.
- K. F. Tipton. *Enzyme Assays*, chapter 1, pages 1–44. Practical Approach. Oxford University Press, Oxford, England, 2nd edition, 2002.
- G. Wahba. *Spline Models for Observational Data*, volume 59 of *CBMS-NSF Regional Conference series in applied mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1990.
- K.-P. Zauner and M. Conrad. Enzymatic computing. *Biotechnol. Prog.*, 17:553–559, 2001.
- R. Zembowicz and J. M. Żytkow. Automated discovery of empirical equations from data. In *ISMIS '91: Proceedings of the 6th International Symposium on Methodologies for Intelligent Systems*, pages 429–440, 1991.

J. Żytkow, M. Zhu, and A. Hussam. Automated discovery in a chemistry laboratory. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 889–894, Boston, MA, 1990. AAAI Press / MIT Press.