Exploration and Exploitation in an Artificial Experimenter

Chris Lovell CJL3@ECS.SOTON.AC.UK
School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

Klaus-Peter Zauner KPZ@ECS.SOTON.AC.UK
School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

Steve Gunn SRG@ECS.SOTON.AC.UK
School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

Abstract

An artificial experimenter is a computational implementation of the decision making processes a laboratory experimenter will make. Artificial experimenter's analyse the available data, propose hypotheses to represent the behaviours investigated and design experiments to evaluate or improve those hypotheses. In doing so they perform active discovery. A key problem faced is deciding when to perform experiments that exploit the information held within the current hypotheses to evaluate them and when to perform experiments that explore the parameter space to discover features of the behaviour being investigated not yet identified. As resources in physical experimentation are extremely limited, addressing this trade-off is critical to obtaining a representative model of the system under investigation. To achieve this, a Bayesian notion of surprise has been used to effectively manage the transition between exploration and exploitation in simulated and physical experimental trials.

1. Introduction

In physical experimentation, the resources typically available are generally small in comparison to the size and scale of the parameter space. For example there may perhaps be only a handful of experiments available per parameter dimension. To aid the experimenter, statistical machine learning techniques can be employed to perform pattern analysis on the data available and choose the experiments to perform, with the goal of maximising the information gained whilst minimising the resources spent. These techniques are similar to computational scientific discovery (Langley et al., 1987) and active learning (Cohn et al., 1994; Settles, 2009). We label the combination of techniques that are implemented to perform this pattern recognition and adaptive experiment selection within a laboratory problem, as an artificial experimenter (Lovell et al., 2010). When combined with automated hardware capable of performing the experiments requested, an autonomous experimentation machine can be created as illustrated in Figure 1.

Additionally to the limited resources problem, physical experimentation has the problem that experimental errors or unexpected undetectable physical changes in the reactants can occur, which can yield observations not representative of the behaviour being investigated. These erroneous observations can be thought of as being outliers, except that there will be insufficient data available to identify them as such with any degree of confidence. One approach to handling this uncertainty in the observation validity, is to consider multiple hypotheses in parallel that have different views about the data (Lovell et al., 2010). The information within these hypotheses can be exploited to select experiments where they most disagree, in order to obtain experiments that can dismiss invalid hypotheses. However experiments must also be performed that can search for features of the behaviour that have not yet been characterised. The artificial experimenter has to make a continual trade-off between performing experiments to discover features of the behaviour not yet identified, with spending resources to ensure that the current models of the behaviour being investigated are accurate.

Appearing in *Proceedings of the 28th International Conference on Machine Learning*, Bellevue, WA, USA, 2011.

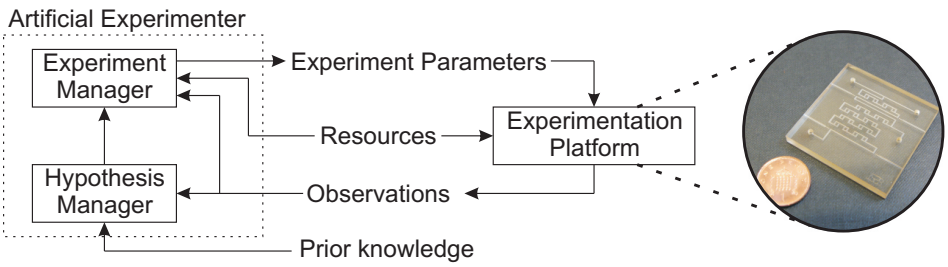

Figure 1. Illustration of a closed-loop autonomous experimentation machine. The artificial experimenter has components to model the data and choose experiments, which are used to provide experiment parameters to the experimentation platform. The experimentation platform then performs the requested experiments, returning the observational measurements as input back to the artificial experimenter. All occurs without human interaction.

Here we discuss how the trade-off between exploration and exploitation has been considered within an artificial experimenter that can perform automatic response characterisation with a small, noisy and potentially erroneous set of observations. In particular enzymatic responses has been considered as the domain to evaluate the techniques, however the design of the artificial experimenter is domain independent. In Section 2 we briefly discuss how hypotheses are represented within the system. In Section 3 we discuss different methods that have been used for managing the exploration–exploitation trade-off within our artificial experimenter and present results showing the performance of those techniques in Section 4. In Section 5 we discuss how previous systems have addressed the exploration–exploitation trade-off.

2. Hypothesis Manager

The hypothesis manager is the component of an artificial experimenter that analyses the available data to form hypotheses about the behaviour being investigated. First we discuss how a hypothesis is represented. A hypothesis provides a model of the behaviour being investigated and allows predictions to be made about regions of the parameter space that have not yet been experimentally investigated. In response characterisation, a suitable method of representing a hypothesis is through a regression technique. Here a hypothesis is represented as a smoothing spline in 1-dimensional parameter spaces, and as a similarly weighted thin-plate spline in higher dimensional parameter spaces.

To deal with the uncertainty caused by only having a limited number of noisy and potentially erroneous observations, a hypothesis manager may consider many hypotheses in parallel (Lovell et al., 2010). Each hypothesis can maintain a different view of the behaviour being investigated, along with different views about the validity of observations. In our design, hypotheses go through a process of refinement in cases where observations do not agree with hypotheses, to develop hypotheses that should be more representative of the true underlying behaviour. An observation and hypothesis are identified as being in disagreement when the observation falls outside of the 95% error bar interval for the prediction of the hypothesis. When refining a hypothesis under these circumstances, the system must take into consideration the problem that it will not know whether the disagreement between observation and hypothesis is because the hypothesis is incorrect, or if the observation is erroneous. The refinement process handles this consideration by creating two new refinements of the original hypothesis that the observation disagreed with. In one refinement the disagreeing observation will be declared to be erroneous and the observation will receive a weight of 0, with all other parameters remaining the same. In the other refinement, the observation will be declared to be valid and the observation will be given a high weight, with all other parameters remaining the same. The zero weight will cause the regression to ignore the observation, whilst a high weight will draw the output of the regression curve closer to the observation. The original and two refined hypotheses are all kept in consideration by the hypothesis manager within a working set of hypotheses.

After each experiment is performed, the hypothesis manager creates a new set of hypotheses with random initial parameters to give different starting views of the behaviour being investigated. These hypotheses are added to the working set of hypotheses that were kept in consideration in previous rounds of experimentation. All hypotheses are then compared to all observations to identify any disagreements, where refinements are made to the hypotheses in cases where there are

disagreements. Finally all hypotheses are evaluated against all of the previous observations to determine their confidence and quality. By maintaining a working set, or committee of different hypotheses, decisions about the shape of the response or validity of observations can be postponed until sufficient data is available to reject incorrect assumptions.

3. Experiment Manager

The experiment manager determines the experiments to perform using the information available within the hypotheses along with information about where in the parameter space previous experiments have been performed. In choosing experiments, the experiment manager needs to ensure that the parameter space is explored to allow for the discovery of new features of the behaviour being investigated, whilst also making sure that data is obtained that can differentiate between the different hypotheses under consideration to identify the most likely candidate. Experiments that explore the parameter space should be placed in regions where there are currently no observations available. Whilst experiments that obtain data to evaluate the hypotheses, should exploit the information held within the hypotheses to determine experiments that differentiate between as many hypotheses as possible per experiment.

Before we consider the trade-off between exploration and exploitation, we briefly define how hypotheses can be efficiently differentiated between. To differentiate between hypotheses, the experiment should be chosen where there is the most disagreement between the predictions of the hypotheses. At first glance it may appear that taking the variance of hypothesis predictions would provide the suitable measure of disagreement. However variance can be made to be artificially large in the presence of a single outlying hypothesis prediction, which can cause experiments to be chosen that only differentiate between the outlying hypothesis and all other hypotheses under consideration (Lovell et al., 2010). Instead we use a maximum disagreement measure that places experiments where there are differences between high quality hypotheses that currently agree on the previous observations obtained, defined in (Lovell et al., 2010):

$$D(x) = \sum_{i=1}^{|\mathcal{H}|} \sum_{j=1}^{|\mathcal{H}|} \left(1 - P_{h_i}\left(\hat{h}_j(x)|x\right)\right) Q(h_i, h_j) \quad (1)$$

where $\mathcal{H}$ is the set of working hypotheses under consideration, P_{h_i} is the probability that hypothesis h_i agrees with the prediction of hypothesis h_j for experiment parameter x, defined as:

$$P_{h_i}\left(\hat{h}_j(x)|x\right) = \exp\left(\frac{-\left(\hat{h}_i(x) - \hat{h}_j(x)\right)^2}{2\sigma_i^2}\right) \quad (2)$$

where $\hat{h}$(x) is the prediction of a hypothesis for x, σ_i is the error bar at x for hypothesis h_i. The term $Q(h_i, h_j)$ is the measure of quality and agreement between hypotheses, defined as:

$$Q(h_i, h_j) = C(h_i)C(h_j)A(h_i, h_j) \quad (3)$$

where $C(h_i)$ is the confidence of hypothesis h_i based on the previous N observations, defined as:

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

with τ kept constant at 1.96. The function $A(h_i, h_j)$ is the agreement between the hypotheses for the previous observations obtained, defined as:

$$A(h_i, h_j) = \frac{1}{N}\sum_{n=1}^{N} \exp\left(\frac{-\left(\hat{h}_i(x_n) - \hat{h}_j(x_n)\right)^2}{2\sigma_i^2}\right) \quad (5)$$

with σ_i again being the error bar value of hypothesis h_i for experiment parameter x. The value of $D(x)$ will be high where there are confident hypotheses, which agree on the current available data, but disagree on the outcome of the proposed experiment. By performing an experiment where $D(x)$ is maximal, evidence should be obtained to identify faults within currently well performing hypotheses that have been identified by other hypotheses.

Next we consider the trade-off between experiments that explore the parameter space and experiments that exploit the information within the hypotheses. In the following we consider two techniques for experiment selection. The multiple peaks technique attempts to select experiments that have a combined ability to explore and exploit. Whilst the surprise technique evaluates how surprising the last experiment obtained was, to determine whether the next experiment should be an exploration or exploitation experiment. In both strategies, the artificial experimenter requires a number of exploratory experiments that can be used to generate an initial set of hypotheses. In the 1-dimensional case presented here the technique performs 5 initial experiments that are equally spaced across the parameter space.

3.1. Multiple Peaks

The exploitation function $D(x)$ in Equation 1, will give a maximal value where the hypotheses most disagree. Performing these experiments when there are good hypotheses in consideration, will identify the hypothesis that most suitably describes the underlying behaviour by disproving the alternate hypotheses. However these exploitation experiments will likely focus on particular areas of the parameter space and may place experiments close to each other in the parameter space. This will mean that little exploration will occur and features of the behaviour may be missed, or only a small number of the differences between hypotheses are examined.

Instead if we consider $D(x)$ across the parameter space, we may expect to see local maxima, or peaks, in the function. Each of these peaks should provide an area of the parameter space where the hypotheses disagree, potentially for different features of the behaviour. If experiments are placed at these peaks, there are three potential benefits. First, there will be a guaranteed information gain through identifying a difference between the hypotheses. Second, different differences between hypotheses will be examined. Finally, experiments will be placed across the parameter space allowing for some additional exploration.

The process of this technique is as follows. After building the initial set of hypotheses, a set of experiments are then chosen that are at the peaks of the discrepancy equation $D(x)$. These experiments are performed sequentially, with the hypotheses updated after each experiment is performed. When all experiments in the set have been performed, the discrepancy equation is recalculated and the next set of experiments are selected and performed.

3.2. Surprise

Several previous artificial experimenter techniques discuss the notion of surprise in scientific discovery, and have employed different formulations of surprise to base their experiment selection strategy (Kulkarni & Simon, 1990; Pfaffmann & Zauner, 2001). These techniques are explored further in the related work in Section 5. Surprising observations are important, as they signify that an outcome occurred that was not expected. It could be that the observation was surprising because it was erroneous, which would require investigation to identify the error and remove it from consideration in the hypotheses. Alternatively an observation could be surprising because the current hypotheses are invalid for the behaviour being investigated. In this instance, further investigation should be made in the region of the parameter space where the surprising observation was found, to allow for more representative hypotheses to be made. Regardless of the cause of a surprising observation, further experiments should be performed when a surprising observation is obtained, to investigate why the observation was surprising.

A Bayesian formulation for surprise exists within the background literature, which uses a Kullback-Leibler divergence to identify surprising differences between prior and posterior hypotheses to identify surprising occurrences in video sequences (Itti & Baldi, 2009):

$$S_o = \int_{\mathcal{H}} P\left(h|D\right) \log \frac{P(h|D)}{P(h)} dH \tag{6}$$

where $P(h|D)$ is the posterior probability for the hypothesis and $P(h)$ is the prior probability. However, for use in an artificial experimenter, this surprise function requires an adjustment. In the current form, the equation identifies surprising improvements to the posterior model and scales the result by the posterior confidence. In an experimental setting, as hypotheses can only be disproved (Chamberlin, 1890), we are more interested in those observations that provide evidence that reduce the confidence in previously good hypotheses. In other words, we are interested in observations that disagree with the hypotheses that are currently viewed as being the most accurate representations of the underlying behaviour under investigation. To make this adjustment, we swap the prior and posterior terms in the function (Lovell et al., 2011). Although it may appear counter intuitive to prefer experiments that weaken the confidence of hypotheses, by identifying the inaccuracies of a hypothesis, the hypothesis will subsequently be refined into a new hypothesis that is more representative of the true underlying behaviour being investigated.

We use this metric to calculate the surprise of the most recently obtained observation. To calculate surprise, we take the confidence of the hypotheses before an experiment is performed to provide the prior probability. The posterior probability is taken as the confidence of the same set of hypotheses directly after the experiment has been performed, where the set of observations to evaluate with will now include the observation obtained in that experiment. This allows the surprise of the most recent experiment across the set of working hypotheses under consideration to be calculated as:

$$S = \sum_{i}^{|\mathcal{H}|} C(h_i) \log \frac{C(h_i)}{C'(h_i)} \tag{7}$$

From this measure of surprise, the decision about whether to explore or exploit next can be made by

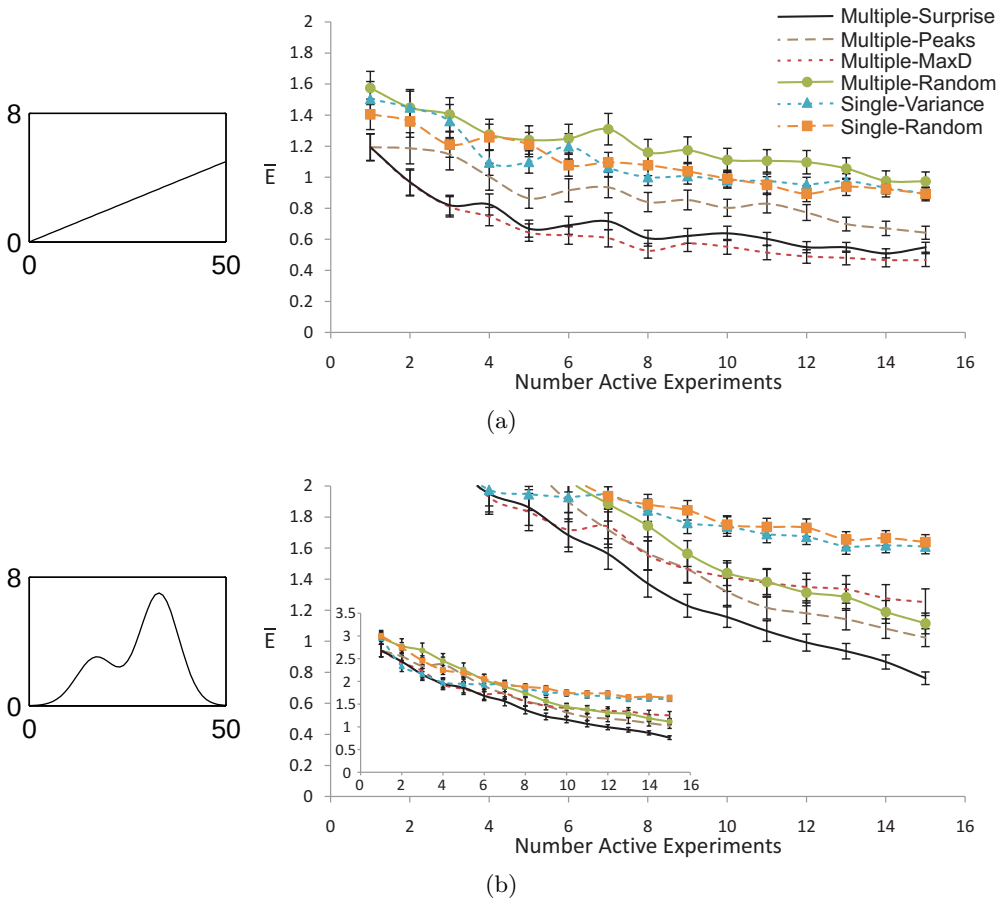

Figure 2. Error between most confident hypothesis and true underlying behaviour over number of actively chosen experiments, for different experiment selection strategies. On the left the underlying behaviour is shown. On the right the error is shown over 100 runs for each technique.

exploiting when the last observation is surprising, and exploring when the last observation was not surprising, as defined as:

$$x = \begin{cases} D(x) & \text{if } S > 0 \\ E(x) & \text{otherwise} \end{cases} \tag{8}$$

where $E(x)$ is a method for choosing an exploration experiment, for example the maximum distance in the experiment parameter space from any previously performed experiment.

The reasons for this trade-off are that when an observation is obtained that is not surprising, so agrees with the current hypotheses, we can infer that the confident hypotheses under consideration agree and a good representation of the behaviour for the features discovered has been made. If a good representation exists for the data available, then features of the behaviour not yet discovered should now be sought after through exploration. Whilst when an observation is surprising, the hypothesis manager will ensure that there are hypotheses that will have opposing views of the surprising observation, meaning that an exploitation experiment can be performed to identify the hypotheses that make the correct assumption about the surprising observation. It may be that several exploitation experiments are performed in a row that investigate one particular feature repeatedly, to allow for refinements of the hypotheses to be made. Once the most confident hypotheses provide a representation that describes that feature well, the observations will become unsurprising again and exploration will occur.

The process of experiment selection occurs as follows. On the final experiment of the initial experiments and all subsequent experiments, the surprise of the observation obtained is calculated. After the surprise of the observation has been calculated, the hypothesis manager updates and refines the hypotheses using the process described previously. If the surprise is positive, then then next experiment will be an exploitation experiment chosen as the maximum of $D(x)$. If the sur-

prise is not positive, then the next experiment is an exploration experiment, chosen as the experiment that is maximally away from all other previously performed experiments in the experiment parameter space.

4. Results

A summary of results are given here, further analysis can be found in (Lovell et al., 2011). In a simulated trial, underlying behaviours representative of potential enzymatic responses were used, represented as a function in the form:

$$y = f(x) + \epsilon + \phi \tag{9}$$

where $f(x)$ represents the behaviour, ϵ represents Gaussian noise applied to each observation with $N(0, 0.5^2)$ and ϕ represents the occurrence of noisy results by applying shock noise of $N(3, 1)$ to three randomly selected experiments. Each run began with 5 initial exploratory experiments, followed by 15 actively selected experiments using a particular experiment selection strategy. The experiment selection strategies were: random selection, the maximum of the discrepancy equation $D(x)$, multiple peaks, and surprise based exploration–exploitation switching. Also examined were techniques employing only a single hypothesis, where experiments were selected through random selection and where the error bars of the hypothesis were maximal. Error was measured as the mean squared error between the most confident hypothesis and the true underlying behaviour after each experiment. Two of the behaviours used are shown in Figure 2.

In simpler monotonic behaviours, the performance of the surprise technique for exploration–exploitation switching is similar to random experiment selection and better than the multiple peaks technique, as shown in the example given in Figure 2(a). However, in more complex non-monotonic behaviours, the surprise technique outperforms all other strategies, as shown in the example given in Figure 2(b). In all behaviours trialled in (Lovell et al., 2011) and in a laboratory trial discussed in (Lovell et al., 2010), the surprise technique performs consistently well. The surprise technique performs well because it is 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. Whilst the discrepancy peaks performs poorly as often the internal trade-off between choosing experiments that have an ability to explore and exploit, lead to experiments being chosen that do neither particularly well. In the discrepancy peaks technique the experiment selection strategy often places repeated experiments within a small area of the parameter space without exploring elsewhere.

Throughout, the surprise technique chooses to perform only a small number of exploration experiments per run. Out of the 15 experiments that the techniques could choose per run, the surprise technique would choose an average of 5 exploration experiments in the simpler monotonic behaviours and 4 in the non-monotonic behaviours. In the monotonic behaviours, the additional exploration experiments were generally chosen in the latter stages of experimentation, caused by the hypotheses agreeing on the behaviour and the technique searching for new behaviours. In the non-monotonic cases the additional exploration experiments were generally chosen in the middle to latter stages of experimentation, which allowed the technique to search for features of the behaviour not yet captured by the hypotheses. In results currently under review, the surprise technique continues to outperform the random and multiple peaks techniques for behaviours with a 2-dimensional parameter space.

5. Related Work

The exploration vs. exploitation trade-off, although not always specifically mentioned, has been addressed by many of the computational scientific discovery systems in the literature. In this section we provide a brief review of the techniques employed in these systems to address the exploration–exploitation trade-off.

The KEKADA system was one of the first examples of an artificial experimenter, which formed more structural hypotheses about the mechanisms of reactions (Kulkarni & Simon, 1990). The system attempted to follow the heuristics used by Hans Krebs to discover the urea cycle, to determine whether a computational system could rediscover knowledge obtained in physical experimentation. A key part of the KEKADA system was reacting to surprising observations, which are those observations that do not agree with the current view of the behaviour being investigated. When a surprising observation was obtained, the system would follow several strategies for reacting to the surprise that would exploit the information available. Examples of the strategies employed are: identifying the independent parameter that caused the surprise and identifying errors in the current hypotheses. The system was able to provide a good model of the heuristics used by Hans Krebs, however it could be easily outperformed by human experimenters who had far more heuristics available to them.

The FAHRENHEIT system, an extension of BACON (Langley et al., 1987), was designed to find an empirical theory that could describe an observed behaviour within a parameter space (Żytkow et al., 1990). FAHRENHEIT was demonstrated to work in an autonomous experimentation machine in the field of electro-chemistry (Żytkow et al., 1990). To begin the system performed experiments that explored the parameter space, to allow it to produce models of regularities in the behaviour. When the system discovered an irregularity, for example a rapid change in phase such as a discontinuity, the system would then focus experiments on investigating the extent of the irregularity, producing separate models for the regions of regularity adjacent to the irregularity. However this system would require a large number of experiments to be performed to provide the data required.

A system developed to automate a chemistry workstation, employed a grid search with decreasing grid size as part of its strategy to manage the exploration–exploitation trade-off (Dixon et al., 2002). The goal of the system was to discover the parameters that produced the highest yield within the experiment parameter space. Initially experiments were placed spread out across the parameter space to provide exploration by a grid with large grid squares. The size of the grid squares decreased over subsequent experiments, to provide a more detailed analysis of the behaviour. Additionally a simplex based experiment selection technique was employed in later stages of experimentation, where experiments would focus towards areas of the parameter space where previously high yields were obtained (Du & Lindsey, 2002). Like many evolutionary algorithms, the technique had the potential for becoming stuck within a local maxima.

Scouting was an evolutionary algorithm that evolved parameters based on an adaptive measure of surprise, which was able to manage the trade-off between exploration and exploitation (Pfaffmann & Zauner, 2001). Like KEKADA, surprising observations were those that differed from the hypothesis under consideration. When no observations were surprising, experiments would be placed randomly within the parameter space. When a surprising observation was obtained, the evolutionary algorithm would then place experiments near the surprising observation. Importantly, as more experiments were placed near the initially surprising observation, so the model would better represent the behaviour and the observation would become less surprising. This adaptivity of surprise, meant that once sufficient information had been obtained to investigate why the observation was surprising, the algorithm would again place experiments in other areas of the parameter space, automatically addressing the exploration-exploitation trade-off. The scouting approach was demonstrated within an autonomous experimentation machine to characterise enzymatic response behaviours (Matsumaru et al., 2002). A problem with this technique was that if an observation was erroneous, then it could remain surprising as subsequent observations would not agree with it, meaning that the system could remain performing exploitation experiments in that region.

The robot scientist (King et al., 2004) does not algorithmically consider the trade-off between exploration and exploitation. Instead the system is provided with a large body of information within a relatively small domain, which is then used to formulate hypotheses that can be tested to determine their validity. Essentially the initial information provided by the user is the only exploration that occurs. The active learning technique the system uses to select experiments to examine the hypotheses is purely exploitative, by choosing experiments that will reduce the likely cost to determine the most representative hypothesis.

Another technique that investigates automatically characterising enzymatic response characterisation, performs a largely exploration focussed experiment selection (Bonowski et al., 2010). Initially the technique is explorative through placing experiments using a space fitting algorithm to ensure a good distribution of experiments across the experiment parameter space. Later experiments combine exploitation and exploration through placing experiments where the uncertainty in the hypothesis is greatest, but also explorative through requiring experiments to fulfil a minimum distance requirement between experiments in the parameter space. The minimum distance required decreases over time, to allow finer examination of the parameter space. However the technique does not consider erroneous observations.

6. Conclusion

A key problem faced in automating the process of discovery is managing the trade-off between exploration and exploitation. In active learning scenarios where there is little data to build hypotheses from and where that data may be erroneous, a method of responding to the surprise of an observation provides a good method of managing this trade-off. When an observation is surprising, experiments should be made that identify why the observation is surprising, through exploiting the information held in the hypotheses. When observations are not surprising, then experiments should be performed that search for new surprising features of

the behaviour. By having separate methods of exploring or exploiting the parameter space that are chosen based on the surprise of the most recent observation, a technique has been developed that consistently performs well in automated characterisation with small, noisy and potentially erroneous observations. Whilst we have found that attempting to ensure each experiment has an element of exploration and exploitation, leads to a technique that can perform poorly by choosing experiments that neither explore or exploit particularly well.

One problem with using the surprise technique is determining the number of initial exploratory experiments to perform. By performing too few, the hypotheses will have little data available meaning that many early experiments are likely to provide surprising observations, leading to more exploitation to ensure each feature is highly accurate before searching for further features. Whilst performing too many initial experiments may waste resources by performing experiments in the regions of the behaviour understood by the system and by doing so reduce the resources available to evaluate the hypotheses.

Acknowledgements

The reported work was supported in part by a Microsoft Research Faculty Fellowship to KPZ. This work was supported in part by the IST Programme of the European Community, under the PASCAL2 Network of Excellence, IST-2007-216886. This publication only reflects the authors' views.

References

Bonowski, F., Kitanovic, A., Ruoff, P., Holzwarth, J., Kitanovic, I., Bui, V. N., Lederer, E., and Wölfl, S. Computer controlled automated assay for comprehensive studies of enzyme kinetic parameters. *PLos ONE*, 5(5):1–10, 2010.

Chamberlin, T. C. The method of multiple working hypotheses. *Science (old series)*, 15:92–96, 1890. Reprinted in: Science, v. 148, p. 754–759, May 1965.

Cohn, D., Atlas, L., and Ladner, R. Improving generalization with active learning. *Machine Learning*, 15:201–221, 1994.

Dixon, J. M., Du, H., Cork, D. G., and Lindsey, J. S. An experimental planner for performing successive focused grid searches with an automated chemistry workstation. *Chemometrics and Intelligent Laboratory Systems*, 62:115–128, 2002.

Du, H. and Lindsey, J. S. An approach for parallel and adaptive screening of discrete compounds followed by reaction optimization using an automated chemistry workstation. *Chemometrics and Intelligent Laboratory Systems*, 62:159–170, 2002.

Itti, L. and Baldi, P. Bayesian surprise attracts human attention. *Vision Research*, 49:1295–1306, 2009.

King, R. D., Whelan, K. E., Jones, F. M., Reiser, P. G. K., Bryant, C. H., Muggleton, S. H., Kell, D. B., and Oliver, S. G. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:247–252, 2004.

Kulkarni, D. and Simon, H. A. Experimentation in machine discovery. In Shrager, J. and Langley, P. (eds.), *Computational Models of Scientific Discovery and Theory Formation*, pp. 255–273. Morgan Kaufmann Publishers, San Mateo, CA, 1990.

Langley, P., Simon, H. A., Bradshaw, G. L., and Żytkow, J. M. *Scientific Discovery: computational explorations of the creative processes*. MIT Press, 1987.

Lovell, C. J., Jones, G., Gunn, S. R., and Zauner, K.-P. An artificial experimenter for enzymatic response characterisation. In *13th International Conference on Discovery Science*, pp. 42–56, Canberra, Australia, 2010.

Lovell, C. J., Jones, G., Gunn, S. R., and Zauner, K.-P. Autonomous experimentation: Active learning for enzyme response characterisation. *JMLR: Workshop and Conference Proceedings*, 16:141–155, 2011.

Matsumaru, N., Colombano, S., and Zauner, K.-P. Scouting enzyme behavior. In Fogel, D. B., El-Sharkawi, M. A., Yao, X., Greenwood, G., Iba, H., Marrow, P., and Shackleton, M. (eds.), *2002 World Congress on Computational Intelligence, May 12-17*, pp. CEC 19–24, Honolulu, Hawaii, 2002. IEEE, Piscataway, NJ.

Pfaffmann, J. O. and Zauner, K.-P. Scouting context-sensitive components. In *The Third NASA/DoD Workshop on Evolvable Hardware—EH-2001*, pp. 14–20, 2001.

Settles, B. Active learning literature survey. Technical report, University of Wisconsin-Madison, 2009.

Żytkow, J. M., Zhu, J., and Hussam, A. Automated discovery in a chemistry laboratory. In *Proceedings of the 8th National Conference on Artificial Intelligence*, pp. 889–894, Boston, MA, 1990. AAAI Press / MIT Press.