Hyperspace Graph Paper: Visualising interactions between search algorithms and landscapes

Janet Wiles[1] and Bradley Tonkes[2]

[1]School of Psychology and School of Information Technology and Electrical Engineering
[2]Institute for Molecular Biosciences and School of Information Technology and Electrical Engineering
The University of Queensland QLD 4072 Australia
{j.wiles, btonkes}@itee.uq.edu.au

Visualising landscapes

Hyperspace graph paper (HSGP) is a technique for visualising surfaces defined over moderate dimensional binary spaces[1]. The key to this technique is in unfolding the hypercube using recursive steps so that the topology of the high dimensional space is reflected in a recursive structure in the two-dimensional unfolding (see Figure 1). In this paper, we describe the original studies that tested the ease of using of this layout technique, and then describe a preliminary case study of its application to understanding the interaction between search algorithms and landscape structures.

Hyperspace graph paper was originally designed as a layout technique to assist users in gaining a working knowledge of multidimensional concepts. Human perception can be viewed as a collection of computational primitives that exploit features such as spatial proximity or optic flow to provide immediate insight into low dimensional data (one to three dimensions). Orthographic projections are a common technique for exploring higher dimensional spaces, but can be limited even for properties of such structures such as distance and proximity. The question arises as to whether humans can gain immediate insight into higher dimensional structures if suitable representations are provided. The recursive layout of hyperspace graph paper was developed as a candidate representation for developing such a tool.

A training program was developed to familiarise subjects with the layout of six and eight dimensional hyperspaces, and a simple task of navigation through the space was designed. Each subject who completed the training session also completed a series of trials on tasks of varying difficulty from two to five dimensions. In a pilot study, ten volunteers were tested on the training program. After one hour of training, five subjects were only just beginning to understand the tasks involved, and did not complete the testing phase. They were all arts students who were not familiar with maths and science concepts. The other five students were computer science and mathematics PhD candidates, and all rapidly learned the spatial structures inherent in the layout.

The training program was refined and the original five successful subjects and an additional fifteen subjects were recruited. They were all computer science or mathematics PhD candidates and were paid for their participation. The testing phase consisted of a series of navigation tasks, in which subjects had to determine the number of dimensions that separated pairs of points on hypercubes, represented on the hyperspace graph paper. Simple heuristics for navigation between dimensions on the hypergraph layout were described, but subjects were free to develop their own intuitions and heuristics as well. Measures were taken of the accuracy and time to determine distances for each pair. Surprisingly, even among this group of highly sophisticated users, performance varied markedly. Half of the subjects made substantial numbers of errors on the distance tasks (not significantly different from chance) and their times were excluded from further analysis. Of the remaining ten subjects, accuracy was very high.

We also measured the relationship between the distance (in terms of number of dimensions between points) and response rates. For the majority of the successful subjects (9/10), there was a clear linear relationship. This result was a natural consequence of the navigation heuristics suggested to the users and hence consistent with expectations. The one subject who had a different pattern of responses was an interesting case. He had perfect performance on the tasks, had the fastest response rates of all subjects and also showed no increase in the time to estimate distances for pairs of points that were further apart.

---

[1] An interactive tool is available at http://www.itee.uq.edu.au/~btonkes/hsgp.html

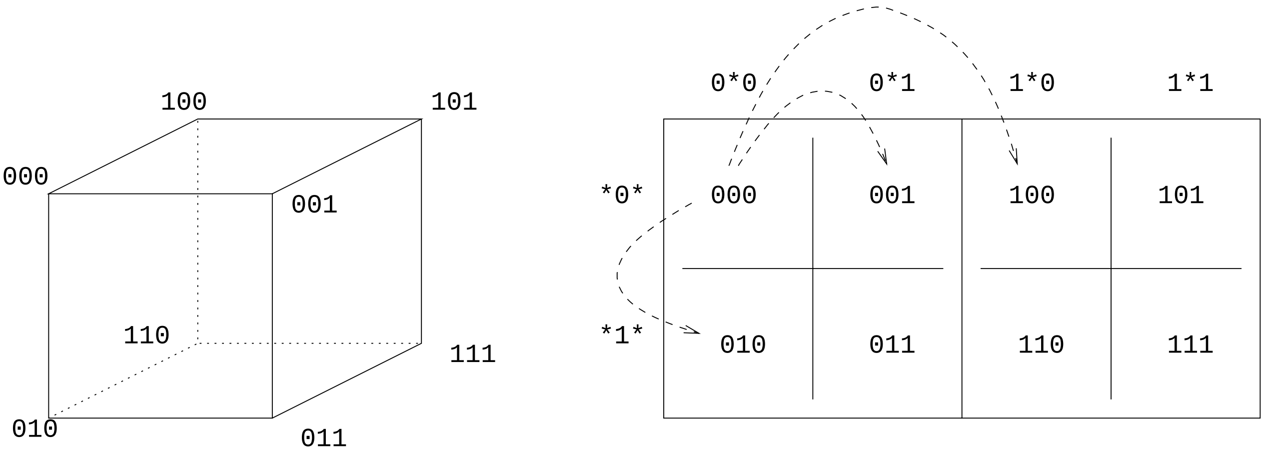

Figure 1. The Hyperspace graph paper layout. Mapping from binary hypercubes to two-dimensions involves a recursive procedure. On the right, the neighbourhood of the point 000 is shown. In general, the neighbourhood structure can be inferred via recursive symmetries. The resulting unfolded graph can be shaded to reflect a fitness function, points of high fitness being coloured in darker shades, as shown in Figure 2.

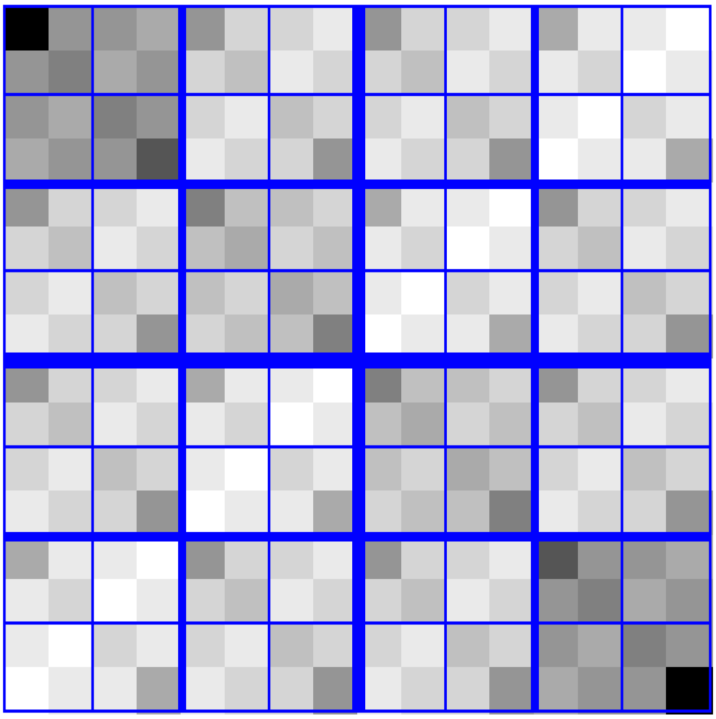

Figure 2. HIFF, the hierarchical-if-and-only-if function, shown over an 8 dimensional space. HIFF is a recursively defined function whose schema comprise strings of either all zeroes or all ones. Thus HIFF has multiple local optima, which in HSGP lie along the top left to bottom right diagonal. HIFF is useful in studying the effectiveness of evolutionary search techniques.

Results

The study showed that the hyperspace graph paper was useful as a high dimensional visualisation for some but not all subjects. The accuracy of each subject was a clear indication of whether the visualisation was a useable tool for that individual. The clear differences in performance showed that some subjects understood and could competently use the layout representation after one hours training, and others were unable to use it effectively at all. This difference reflects the authors' experiences with HSGP, that once user's "get" the layout, they can use it effectively. Prior to that point (which can take some time), the layout seems to have no benefit at all. The linear relationship between distance and response rates for the majority of successful users supports the idea that a heuristic for navigation can be iteratively applied. The one outstanding subject shows that for some people, the tool allows an even more effective insight into higher dimensional spaces.

Following the visualisation tests, a series of tools were developed, with HSGP being the latest and most sophisticated. It allows exploration of fitness landscapes used in evolutionary computation, including a variety of schema-based and multi-modal functions (for details of such studies see Wiles and Tonkes, 2002). HSGP has primarily been applied to visualising surfaces defined in these spaces. Using the technique in this way allows the user to bridge the gap between the limits of the human perceptual system (two or three dimensions) and the properties of higher dimensional surfaces which often cannot be effectively reasoned about with low dimensional metaphors. Because fitness landscapes scale exponentially, there are practical limitations for any visualisation technique that attempts to show an entire space. Hence HSGP has an effective limitation of around 16 dimensions that can be usefully displayed. An understanding of properties of high dimensional search problems provides insight into the types of strategies that optimization algorithms should utilise. Previous studies have considered the visualisation of features of search landscapes.

In this paper we consider the interaction between a search algorithm and the features of a fitness landscape. All search algorithms have an inherent search bias, which can be explored through visualisation even without a detailed understanding of the algorithm. This use of visualisation allows us to tune our intuitions about which algorithm to choose for a given application and why.

One class of optimisation algorithms searches by modelling an underlying probability distribution function (pdf) so that peaks of the landscape correspond to peaks of the pdf. Whereas genetic algorithms maintain a population of good solutions, these algorithms attempt to balance exploitation and exploration by explicitly integrating and maintaining information accumulated from each individual generated. Population based incremental learning (PBIL) is perhaps the most simple algorithm in this class (Baluja, 1994).

Here we present a case study in which PBIL attempts to optimize a multi-modal multi-peaked fitness function, hierarchical-if-and-only-if (HIFF, Watson and Pollack, 1999). HSGP is used to unfold the high dimensional pdf in much the same way as we have done for the fitness function in previous studies (see Figure 2).

HSGP clearly reveals the search mode of PBIL (see Figure 3). As can be seen from Figure 3, on the HIFF function, each trial of PBIL finds a single local optimum. Exploration with the visualisation tool allows repeated trials, and the development of intuitions about PBIL's strategy (compare Figures 3c and 3d). PBIL attempts to model the pdf of a function by representing the probability density as the linear combination of probabilities of each dimension. The limitation of PBIL is that it fails to model the joint probability distribution function that underlies the multi-modal structure of HIFF.

Such visualisations could be extended a wider range of search landscapes to see how PBIL performs on each. Additionally, it could be used to explore a wider variety of search algorithms such as the Bayesian optimisation algorithm, which does model the joint probabilities (Pelikan and Goldberg, 2000). By visually classifying search problems, and seeing how algorithms interact with them, insight can be gained that enables us to choose algorithms for tasks of interest.

Acknowledgements

This research was supported by grants from the School of Psychology and Australian Research Council to JW. Thanks to Cathy Mackie who helped with collection of data and Stephen Wong who developed an early prototype of the HSGP layout.

References

Baluja, S. (1994). Population-Based Incremental Learning: A Method for Integrating Genetic Search Based Function Optimization and Competitive Learning. CMU-CS-94-163, Carnegie-Mellon University.

Pelikan, M. and Goldberg, D.E. (2000). Hierarchical Problem Solving by the Bayesian Optimization Algorithm. IlliGAL Report No 2000002, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign.

Watson, R.A. and Pollack, J.B. (1999). Hierarchically-consistent test problems for genetic algorithms.

Proceedings of 1999 Congress on Evolutionary Computation, eds. P. J. Angeline and Z. Michalewicz and M. Schoenauer and X. Yao and A. Zalzala, IEEE Press, 1406—1413.

Tonkes, B. (2002). Hyperspace graph paper WWW page. http://www.itee.uq.eu.au/~btonkes/hsgp.html

Wiles, J. and Tonkes, B. (2002). Visualisation of hierarchical cost surfaces for evolutionary computation. In *Proceedings of the 2002 Congress on Evolutionary Computation*, 157–162.

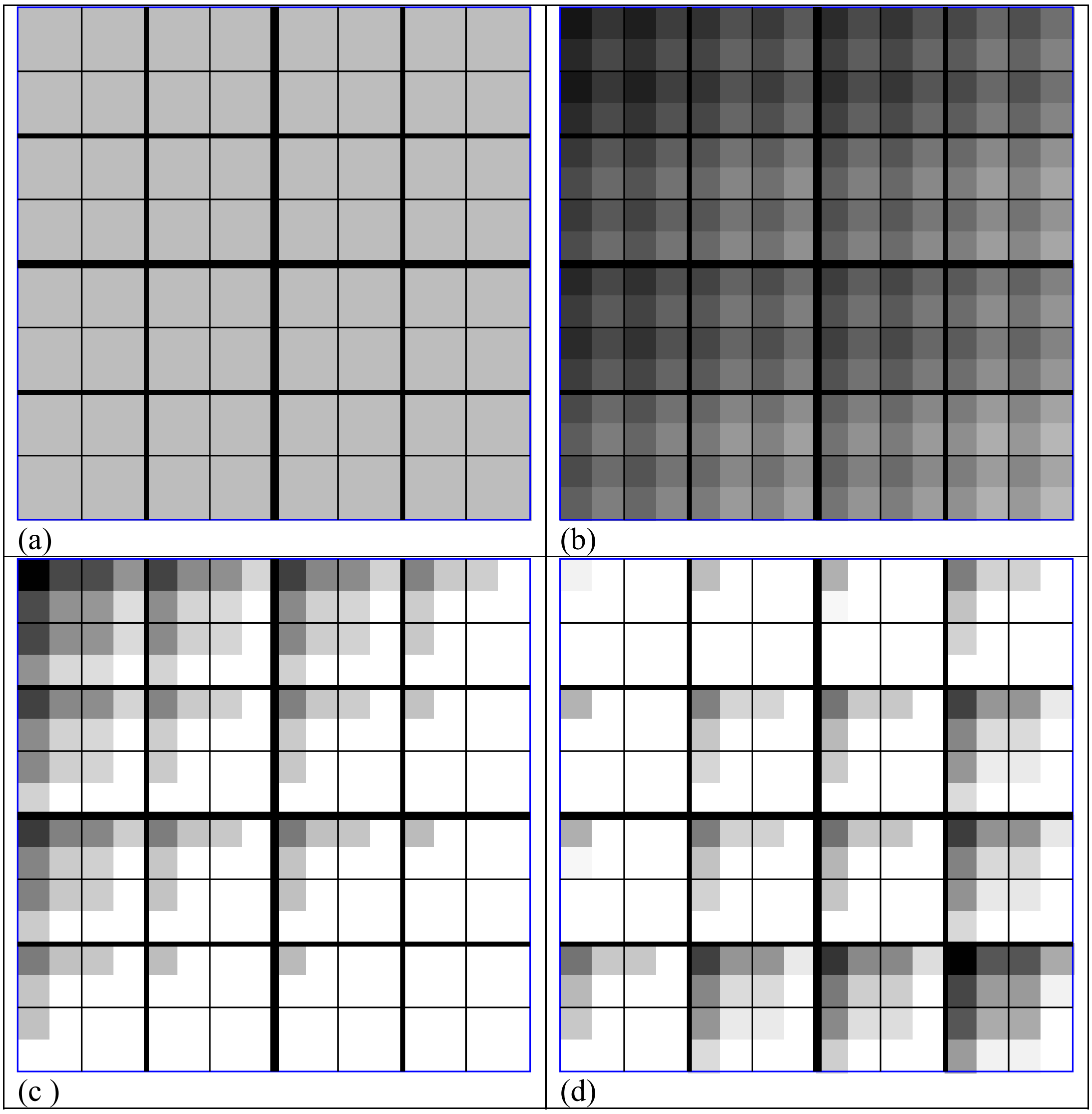

Figure 3. PBIL's progress in optimising HIFF. Each figure shows the probability distribution function (pdf) as constructed by PBIL. (a) The initial uniform distribution, (b-c) As the algorithm progresses, structure begins to emerge so that by the final figure, one of the two global optima (0000 0000) is clearly the most likely solution predicted by PBIL. Note however that as the algorithm progresses the substructure of HIFF is not reflected in the structure of the pdf. (d) An alternative run of PBIL over HIFF shows it converging on a suboptimal solution (1111 1100). As in (c), the substructure of the fitness function is not modelled by the pdf.