Peter Hatfield (Natural Sciences) – Caltech, Summer 2012

I am a current Part III Astrophysics student, who took the Mathematics Tripos for my first three years in Pembroke. Over the summer of 2012, I undertook a research project in the Behavioural Biology department at the California Institute of Technology (Caltech). I worked mainly with postdoctoral staff member Dr Nikhil Joshi.

My project was to investigate community detection in "information networks". Community detection is a commonly studied problem in graph theory that seeks to identify structure in large networks. For example, if one considers the graph of people as nodes connected if the two individuals are friends, the natural communities of schools, cliques of friends' etc. form naturally more densely connected subsets of the network. Discovery of these communities is of great importance and is a matter of great study in social science, genetics, data analysis, neurology and many other fields. The main motivation for studying community structure in the group I was with (Koch Lab) was the application to neurons in the brain to see which operate together as a single mechanism and which are mostly causally disconnected.

The first segment of my project centred on the exploration of a moderately new definition (and algorithm for finding) community structure, namely the partition that maximises the quantity Q, the "modularity", and a spectral matrix method. I developed a couple of implementations of this algorithm, and explored some of the weaknesses of the definition and under what conditions it can fail to spot the "intuitive" partition. Once I understood this definition of community structure well, I started to attempt to apply it to "information networks". The paradigm we were investigating consisted of a system of nodes that could either be "on" or "off". Then at discrete time steps the nodes would update based on their previous states, for example node A might be an AND gate based on what nodes B and C were at the previous time period. The objective was to find a good way of defining community structure in such a network, and find a good way of finding it (the number of partitions of n objects goes as the super exponential Bells number). The initial approach was to convert this updating mechanism into a weighted graph by some means and then use existing community structure techniques to study the system. This methodology gives expected results in strongly differentiated communities, but it is unsatisfactory as there is no a priori best weighted graph to construct, and it is unclear whether this is even the "best" way of looking at a mechanism, for example decomposition (overlapping partition) may be more meaningful. The reason for the added complexity is that mechanisms have more structure than graphs; using the tools of information theory, the relation between two nodes can be simplistically characterised as dependant, synergistic or redundant. Node A can derive information from nodes B and C in 4 essentially different ways; the information can be got from ONLY B, ONLY C, B OR C (redundant information) or ONLY B AND C (synergetic information). It is unclear which should be given the greatest weight in partitioning decisions, especially since the situation gets more complicated with more nodes and once the fact that the output nodes can be grouped by whether they have the same or different information etc.

My first method was to construct an adjacency matrix by taking the entries as the mutual information between node i at the previous time period and node j at the current time. This gives some good results, but is flawed in that it can only spot relations where the output node can derive the information from the input node as a singleton. Critically, for example, it cannot spot XOR relations, because knowing one of the input nodes of an XOR formally gives no information about the output; both must be known. My later development of this technique was to instead use the mutual information *conditional on all the other input nodes*. This essentially measures how much information it is essential to know the input node for. This adjacency matrix gives results consistently in agreement with the intuitive partition and was robust to even rather unusual mechanisms. It remains a "non-a priori" methodology but has already shown use in elucidating community structure in some mechanisms, particularly since it only has $O(n^3)$ run times compared with O(Bells Number), the run time of the previous best method. Towards the end of my project I also developed a definition of community structure that was a priori; essentially finding the dendrogram that minimises "search time" for information, but it was unclear to what extent this agreed with usual concepts of community.