Using Surrogate Models for Shape Optimisation

How can we optimise the design of a machine component when the problem is too complex for traditional methods? In this thesis we explore a surrogate based approach for solving these problems.

Traditionally when designing a new component, an initial design is made and the component's performance is measured. Most likely, the first version does not perform as required, and the design is changed based on the initial tests. When the constructor has designed the second version of the component, its performance is once again tested, and once again redesigned if needed. This loop continues until the performance of the part is satisfactory. This process can be time consuming and can take many iterations before the designed part performs satisfactory. Our thesis will attempt a different approach of designing a component. Instead of designing a component and seeing what its performance is, we do the opposite; we determine a satisfactory performance and use optimisation to find the design that yields that performance. The framework we will implement is based on Bayesian optimisation. Bayesian optimisation consists of three main parts: an initial sampling, a surrogate model, and an acquisition phase.

The first part of Bayesian optimisation is sampling methods. Sampling evenly in one dimension is simple, but in multiple dimensions it is not as simple. Spacing points evenly, while not reusing previous coordinates, becomes more difficult as the amount of dimensions increase. Three different sampling methods, all aiming at spacing points evenly while not reusing coordinates, have been explored and tested throughout this work. Relevant literature state the different sampling methods all have advantages and disadvantages. By varying parameters and testing the performance of the sampling methods, we were able to to verify some of the advantages and disadvantages from literature, while finding others relevant to the case at hand. In the end, two of the methods performed well, while the third underperformed based on the expectations set by literature.

The second, and most important, part of Bayesian Optimisation is surrogate models. A surrogate model is a model, or functional form, that a unknown function is assumed to take. This can be a e.g a polynomial or a Gaussian process. A polynomial is too simple to work in many applications. Instead, a Gaussian process is a good choice since Bayesian Optimisation is built on statistics. Our testing showed a Gaussian process is able to adequately interpolate a function, verifying it is a good choice for a surrogate model.

The last part of Bayesian Optimisation is acquisition functions. They provide a good way improving the surrogate model while lessening the computational cost, when the unknown function is expensive to evaluate. Three different acquisition functions were tested in this thesis. By only valuing high values these functions improve the surrogate model in regions of interest, disregarding regions of low values. One of the functions, despite being the most popular in the scientific community, proved too inflexible and easily got stuck, signifantly increasing the computational cost of the algorithm. The other two had different work arounds, but both yielded similar results, significantly improving the model.

The thesis shows that it is possible to use a surrogate based approach to solve problems too complex for traditional structural optimisation methods. Although, the methodology created is not complete and needs to be tested more thouroughly and tuned to consistently yield good results.