The goal of this thesis was to implement a practical tool for optimizing hyperparameters of neural networks using Bayesian optimization. We show the theoretical foundations of Bayesian optimization, including the necessary mathematical background for Gaussian Process regression, and some extensions to Bayesian optimization. In order to evaluate the performance of Bayesian optimization, we performed multiple real-world experiments with different neural network architectures. In our comparison to a random search, Bayesian optimization usually obtained a higher objective function value, and achieved lower variance in repeated experiments. Furthermore, in three out of four experiments, the hyperparameters discovered by Bayesian optimization outperformed the manually designed ones. We also show how the underlying Gaussian Process regression can be a useful tool for visualizing the effects of each hyperparameter, as well as possible relationships between multiple hyperparameters.