

Gaussian process models are commonly used as emulators for computer experiments. However, developing a Gaussian process emulator can be computationally prohibitive when the number of experimental samples is even moderately large. Local Gaussian process approximation (Gramacy and Apley, 2015) has been proposed as an accurate and computationally feasible emulation alternative. However, constructing local sub- designs specific to predictions at a particular location of interest remains a substantial computational bottleneck to the technique. In this paper, two computationally efficient neighborhood search limiting techniques are proposed, a maximum distance method and a feature approximation method. Two examples demonstrate that the proposed methods indeed save substantial computation while retaining emulation accuracy.