

FocusedILS

The question is how to choose the optimal number of training instances, N ?

- Using too small N leads to good training performance, but poor generalization to previously unseen test benchmarks.
- On the other hand, we cannot evaluate every parameter configuration on an enormous training set - if we did, search progress would be unreasonably slow.

FocusedILS is a variant of ParamILS that **adaptively varies the number of training samples** considered from one parameter configuration to another in order **to focus samples on promising configurations.**

- $N(\theta)$ denotes the number of runs available to estimate the cost statistic $c(\theta)$.

FocusedILS

The question is how to choose the optimal number of training instances, N ?

- Using too small N leads to good training performance, but poor generalization to previously unseen test benchmarks.
- On the other hand, we cannot evaluate every parameter configuration on an enormous training set - if we did, search progress would be unreasonably slow.

FocusedILS is a variant of ParamILS that **adaptively varies the number of training samples** considered from one parameter configuration to another in order **to focus samples on promising configurations.**

- $N(\theta)$ denotes the number of runs available to estimate the cost statistic $c(\theta)$.

The question is how to compare two parameter configurations θ_1 and θ_2 for which $N(\theta_1) \leq N(\theta_2)$?

- *What if we computed the empirical statistics based on the available number of runs for each configuration?*

Can lead to systematic bias if, for example, the first instances are easier than the average ones.

Recommended Material

Frank Hutter, Holger H. Hoos, Kevin Leyton-Brown, and Thomas Stützle: ParamILS: An Automatic Algorithm Configuration Framework. In *Journal of Artificial Intelligence Research (JAIR)*, volume 36, pp. 267-306, October 2009.

Other papers and SW available at <http://www.cs.ubc.ca/labs/beta/Projects/ParamILS/>

