Line/Subspace Search

Dimensionality Reduction Accelerates Optimization/Learning

A useful heuristic is to focus search in a subspace for a time, and then change the subspace.

Line search:

Select a direction (usually the gradient).
Find a minimum along that direction.
Iterate.

Note that in one dimensional search a local minimum can be bounded (for a smooth function).

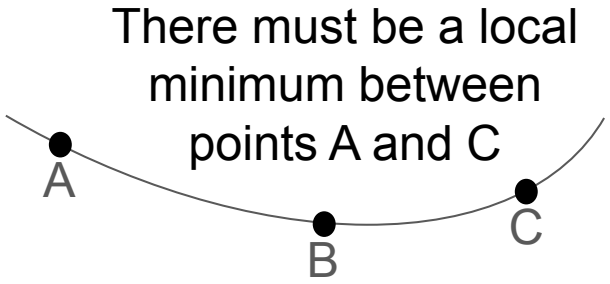

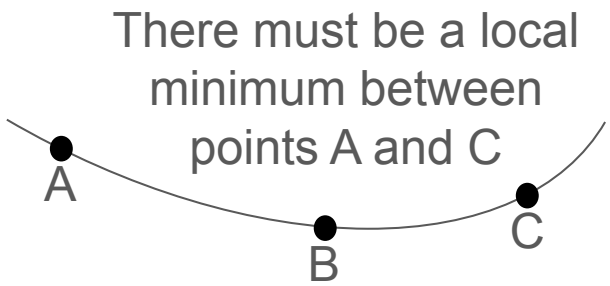
There must be a local minimum between points A and C
A
B
C

A useful heuristic is to focus search in a subspace for a time, and then change the subspace.

Subspace search:

- Choose a subspace
- Optimize within the subspace
- Iterate

This is similar to what we do when practicing sports, where we typically break the behavior into parts, practice the parts separately (drills), and practice the entire behavior (combining the parts).

This is also related to curriculum learning - start with simple tasks, and progressively learn harder tasks.

Transforming the decision variables is another form of dimensionality reduction. So is changing the features or “level” used in the optimization and learning.

Examples: Optimizing actuator commands at each millisecond vs. commanding longer term primitives.

Navigation example:
Where am I going (address)?
What route will I take?
What steady state driving and turning policies will I use?
What muscle activations will I use?
What sensory feedback do I expect?