Standard API for non-adaptive libraries

Adaptive libraries (BLAS, FFT, NN) continuously assembled from the most efficient routines (species) as plugins

Expose all features influencing optimizations

Data set and hardware features ($\vec{f}$)

System state ($\vec{s}$)

Requirements ($\vec{r}$)

Predict most efficient solutions based on features, system state and requirements when enough knowledge Is collected

Learn to trade off speed, size, accuracy, energy, memory usage, costs and other metrics

Minimal set of optimized routines (winning species)

Check behavior ($\vec{b}$)

Detect and record unexpected behavior, numerical instability and performance anomalies

Shared data sets

Most efficient species

Shared most precise models

Shared routines

Most efficient optimizations

Expose to the community to
- **expose more optimizations**
- **add different algorithms**
- **find missing features**
- **improve predictive models**
- **add more data sets**

Repositories with continuously optimized workloads and kernels