SAND2015-9565C

Sandia National Laboratories

# Ultra-high-dimensional optimization for trade space exploration: challenges and lessons learned

Note to reviewers: all results are notional

Matthew Hoffman

Alexander Dessanti

Stephen Henry

Jack Gauthier

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Abstract


- In projects with many conflicting stakeholder objectives, negotiating compromise requires understanding tradeoffs (see Alex' talk for more background)
- We wish to understand tradeoffs for a class of problems with nonlinear constraints and 30+ nonlinear objectives
- For very-many-objective problems, Pareto dominance is not a differentiator; provides no pressure and confounds GAs
- Existing approaches to many-objective optimization focus on converging to a small portion of the Pareto by incentivizing compromise; e.g. aggregation or nearness to a preferred region
- This can severely obfuscate the tradeoffs between solutions
- We need to go back to the drawing board for our problem

# Typical many-objective philosophy

- Focused on improving *convergence toward* (some portion of) the Pareto, for *fewer than 10 objectives*
- "a small improvement in one objective at a large expense to another is unacceptable"

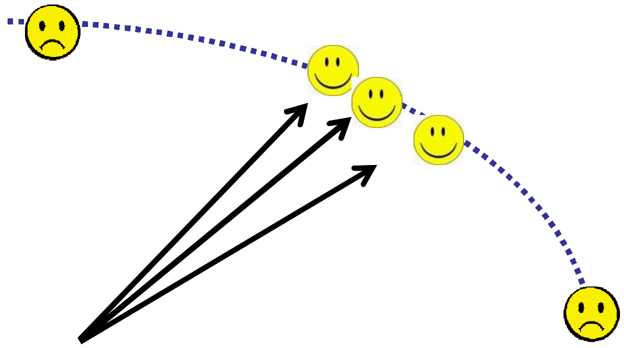

# Our philosophy

- Focused on *characterizing* the Pareto
- Showing *only* compromises gives very little/misleading tradeoff information
- We want to know about the extremes as well as the knee

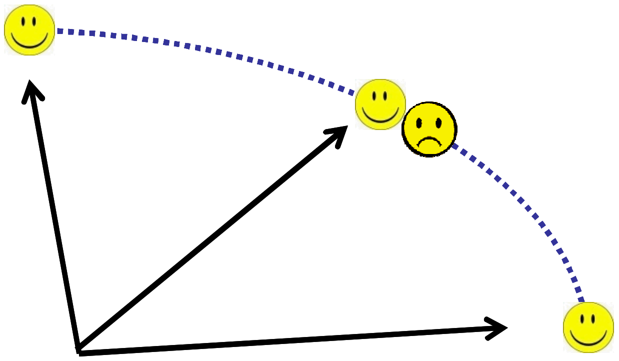

# The problem with *very* many objectives (>10)

- *Everything* is Pareto, so “convergence toward the Pareto” is effectively meaningless
  - One exception: solutions containing genes that are trivially inferior are not Pareto, and it may be very unlikely to dominate/remove these, but these can be detected and removed *a priori* anyway
  - Constraint-infeasible solutions *may* be dominated depending on how infeasibility is penalized, but these can be detected and removed too

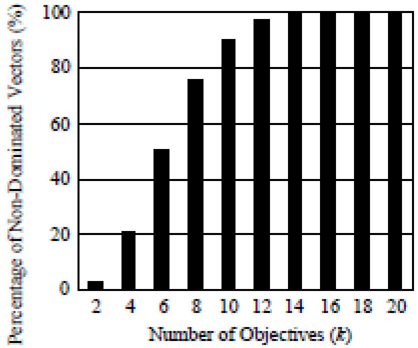


Fig. 1. Average percentage of non-dominated vectors among 200 vectors that are randomly generated in the k-dimensional unit hypercube $[0, 1]^k$.

From H. Ishibuchi, N. Tsukamoto, and Y. Nojima, “Evolutionary many-objective optimization: A short review,” *Proc. of 2008 IEEE Congress on Evolutionary Computation*, pp. 2424-2431, Hong Kong, June 1-6, 2008.

- No dominance-based evolutionary pressure within feasible set
  - If population size unbounded, leads to (effectively) unbounded growth, beyond the limits of system memory or human interpretability
  - If population size bounded, population can be modified and shuffled indefinitely without stabilizing, depending on selection/niching strategy

# How compromises distort tradeoff info

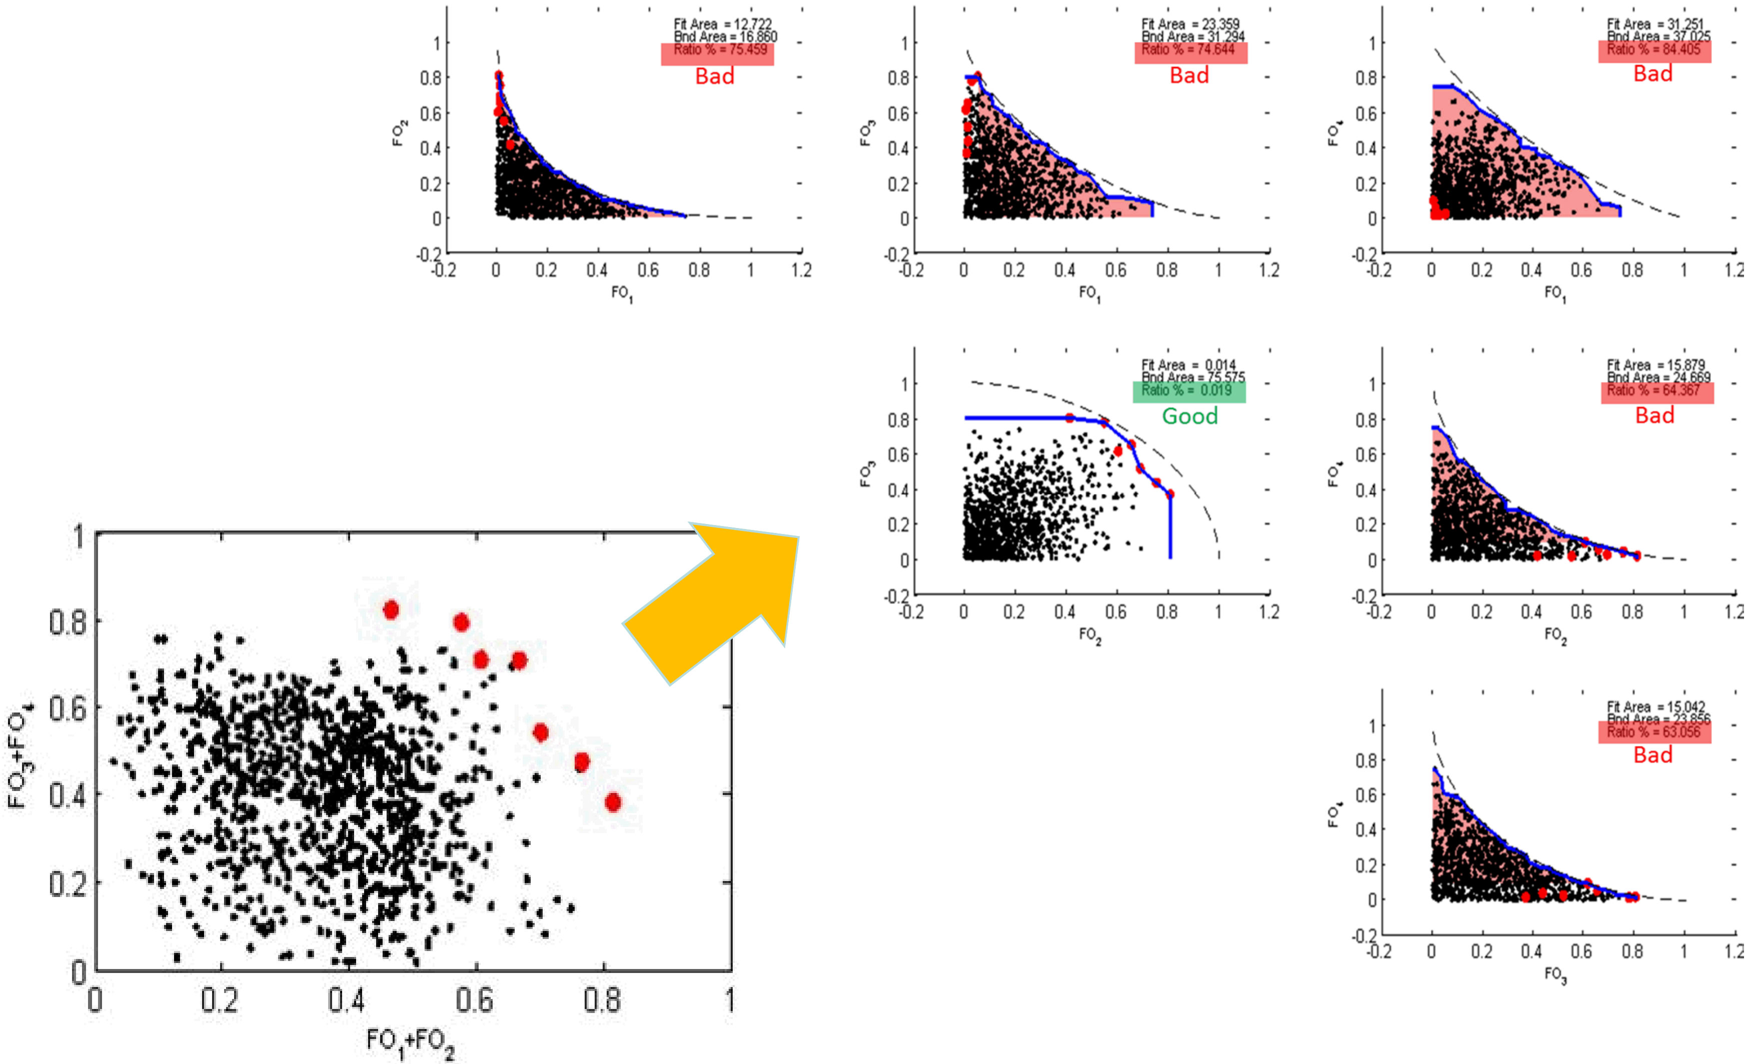

# What we *don't* want

- Loss of information about extreme points
- Clumps (overrepresentation of an area)
- Holes where there shouldn't be (underrepresentation of an area)
- Over-compromising, leading to lack of representation of tradeoffs

# Research Objective

- Our hypothesis is that

a *well-spaced* subset of Pareto points in 30+ dimensions

which includes all 1-d objective optima

will preserve (at least) the most important tradeoff information.

…basically, instead of using the GA to find the “best” set of points, we want to find the most *diverse* subset of *feasible* points.

# Research plan

Sandia National Laboratories

- Proposed process:
  - Find extremal points via 1-d optimization
  - Put these points in the initial population
  - Run GA that
    - preserves extremes
    - preserves diversity
    - stabilizes over time

- Research topics:
  - How to choose which points to preserve when 1-d optima are not unique
  - How to incentivize diversity
  - How to not get stuck with “inbred” families of local optima
  - How to prevent constant population churn
  - How to decide when to stop

# Extreme preservation

Old

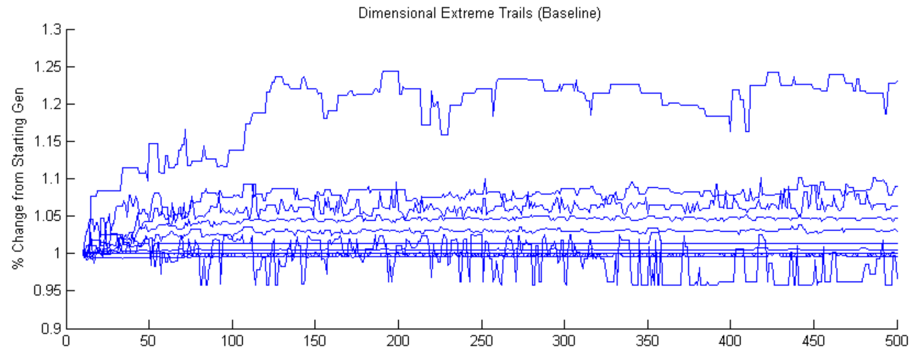


New

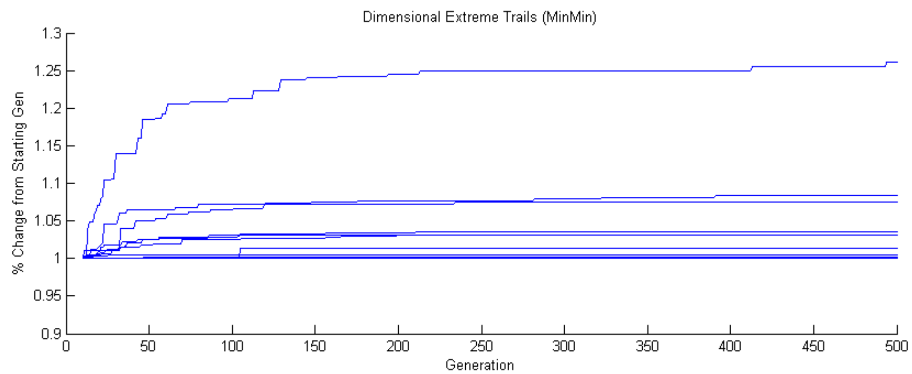


For each dimension $i$, select the optimal point(s) in that dimension. Of these, keep the best for each dimension $j \neq i$.

This will use at most $n^2$ points where $n$ is the number of dimensions.

# Problems with “inbreeding”

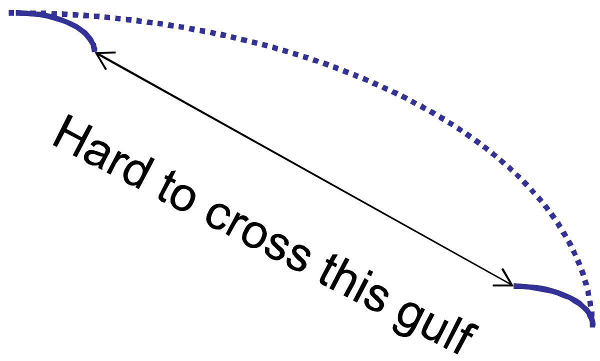


- Mutation unlikely to introduce sufficient genetic diversity
- Crossover problematic
  - Children likely to be dominated by other population members (not a problem in our case)
  - Children likely to be infeasible to
    - Structural constraints (A cannot work with B)
    - Attribute constraints (weight of solution cannot exceed Y)

# Mitigations to “inbreeding”

- Could try to “seed” knee points via 1-d optimization, but this is a hack and becomes quickly unmanageable if you want to do more than 2-d knees or more than just the single X+Y knee for each pair of dimensions.

- Conceptually easy to heal structural dependencies between genes; this is a partial mitigation that appears to work very well

- Not clear how to heal infeasibilities to nonlinear constraints – ongoing research

# Problems with existing niche operator

- Without dominance as a pressure, and with bad/infeasible points removed, selection within the feasible set comes down to measuring and dis-incentivizing crowding.
- Existing niche operator did not stabilize population and led to undesirable emergent clumping behavior

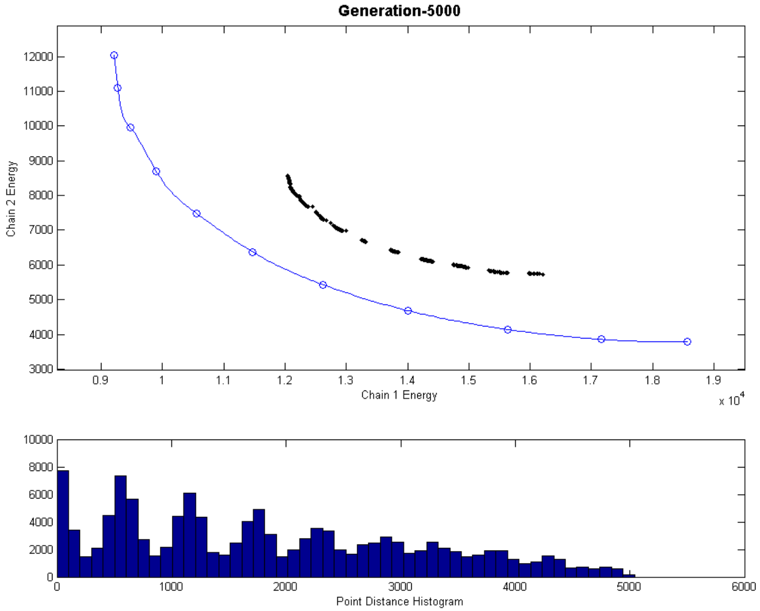


# Niche operator research

- Selection heuristic based on Euclidean distance; tries to greedily maximize minimum distance between selected points
- Works well regardless of original point density

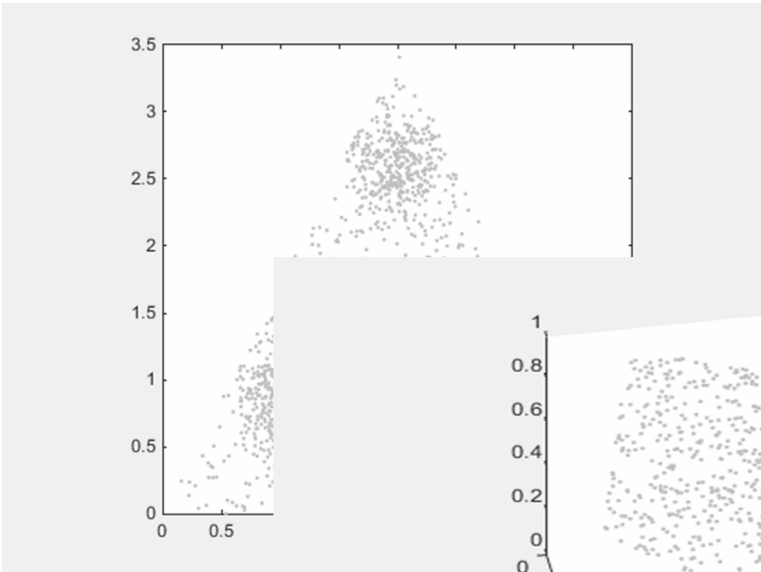


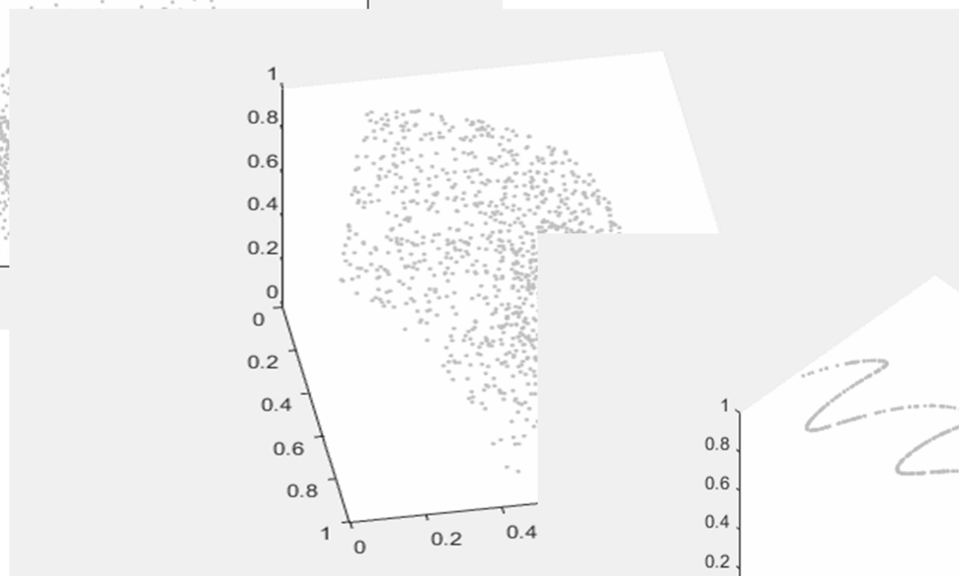


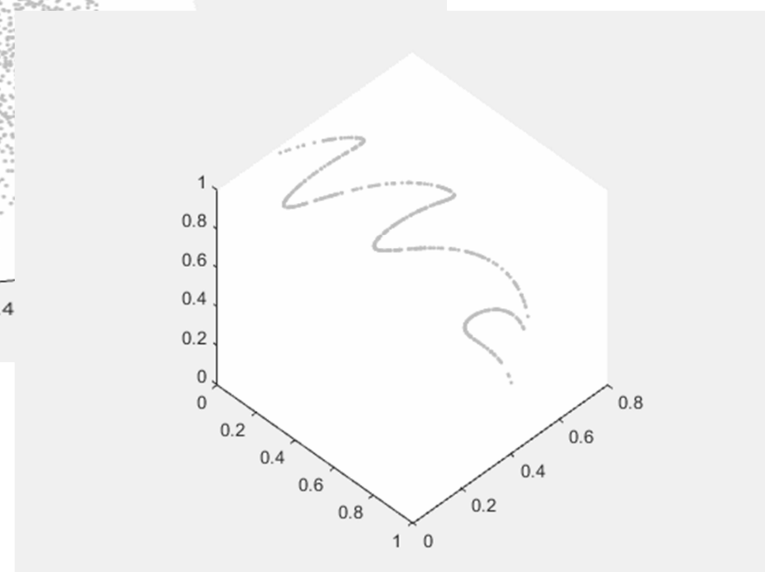


- Works well on sub-dimensional manifolds

# Parameter settings

- Each generation consists of the previous generation and its children – there's no point in having children be identical to parents
- Want every child to be a result of a mutation, crossover, or both

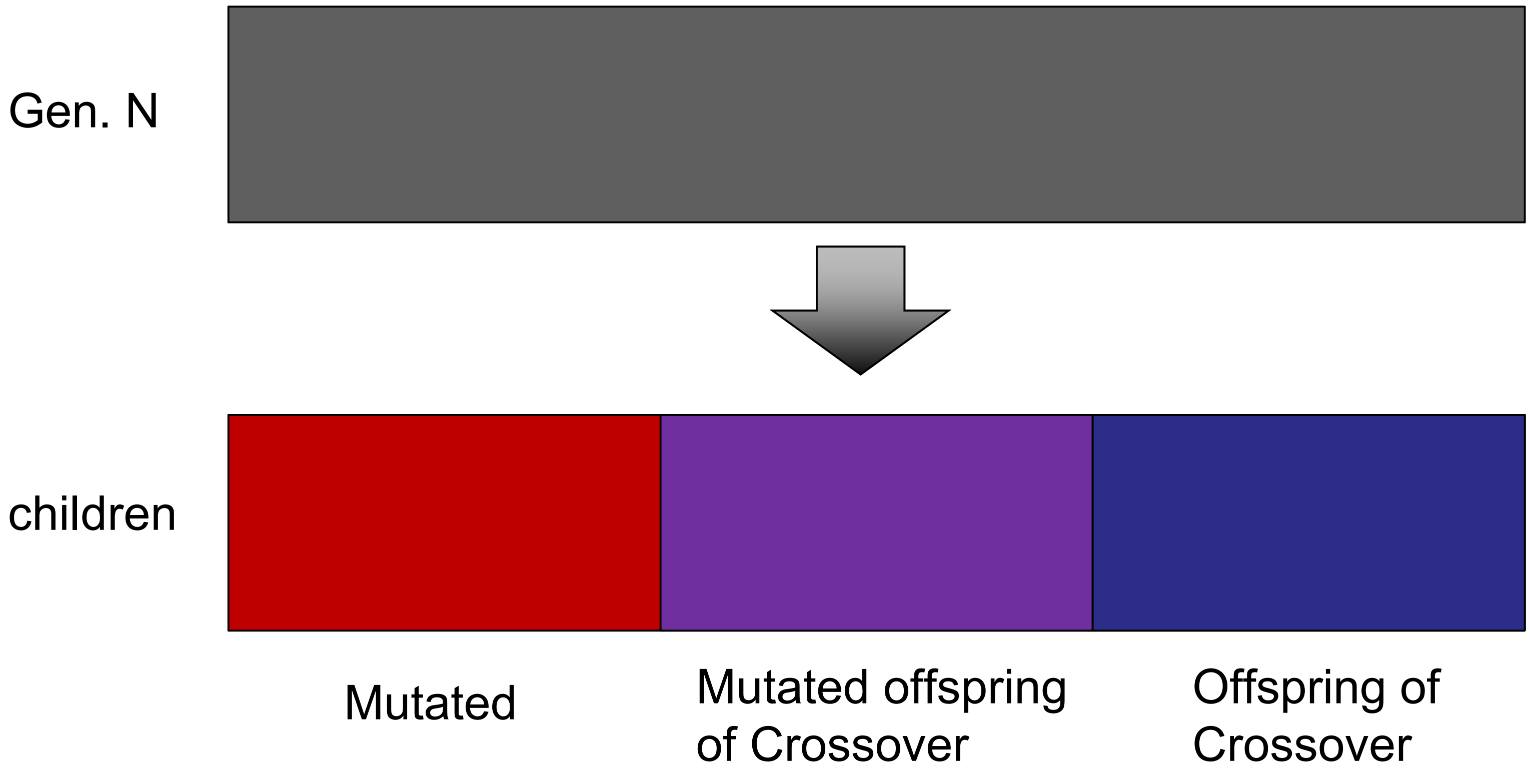

# Conclusion

- (Reiterate abstract)

- Next steps
  - Prove that new niche operator removes emergent clumping issue
  - Determine whether new niche operator also helps with population churn/ nonconvergence issues
  - Experiment with stopping criteria
  - Research ways to make GAs more robust to feasibility dependencies between alleles