Multi-Objective Evolutionary Design of Molecules with Enhanced Nonlinear Optical Properties

Dominic Mashak(1), Jacob Schrum(2), and S.A. Alexander(1)
Department of Physics(1), Department of Math & Computer Science(2), Southwestern University

Introduction

Nonlinear optical (NLO) materials are essential for many photonic, telecommunication, and laser technologies, yet discovering better NLO molecules is computationally challenging due to the vast chemical space and competing objectives. We compare evolutionary algorithms for molecular design with multiple objectives. We encode molecules as SMILES strings and evaluate their properties using quantum-chemical calculations. We compare NSGA-II, MAP-Elites, MOME, a single-objective ($\mu + \lambda$) EA, and simulated annealing. Quality Diversity methods maintain archives across a measure space defined by atom and bond count, enabling the discovery of diverse molecules.

Methods

SMILES (Simplified Molecular Input Line Entry System) strings, encode molecular structures as ASCII strings. We restrict molecules to C, N, O, and H atoms, focusing on organic NLO candidates. Atoms are represented by atomic symbols, but use implicit hydrogens based on valence; bonds limited to single (-)/double (=) bonds; branches are represented by parentheses at the attachment point, and rings by matching digit labels.

We use seven mutation operators to modify SMILES strings:

(1) Change a random bond type (e.g., C=C-N-O → C-C-N-O).

(2) Add a random atom with a random bond (e.g., C=C-N-O → C=C-O-N-O).

(3) Add a random atom as a branch (e.g., C=C-N-O → C=C-N(-O)-O).

(4) Delete a random atom and its bond (e.g., C=C-N-O → C=C-N).

(5) Change a random atom type (e.g., C=C-N-O → C=C-N-C).

(6) Connect two random atoms to form a ring (e.g., C=C-N-O → C1=C-N-O1).

(7) Delete a ring (e.g., C1=C-N-O1 → C=C-N-C).

Quantum-chemical properties are computed with PySCF, a Python library for ab initio electronic structure calculations via Hartree-Fock (HF). We use HF/3-21G.

Experiment

MOO approaches optimize four objectives for electro-optic modulators: the first-to-second hyperpolarizability ratio (β/γ, primary objective), linear polarizability range deviation ($f\alpha$), HOMO-LUMO gap range deviation ($f\Delta E$), and energy per atom ($Etotal/Natoms$):

- First-to-Second Hyperpolarizability Ratio (β/γ): We maximize this ratio to favor strong second-order responses relative to third-order effects, promoting efficient frequency conversion without competing third-order processes.
- Linear Polarizability Range Deviation ($f\alpha$): Large α indicates high electronic mobility, but excessive values cause optical losses/reduced photostability. Targeting 100–500 a.u., we minimize: $f\alpha = \max(0, 100 - \alpha) + \max(0, \alpha - 500)$.
- HOMO-LUMO Gap Range Deviation ($f\Delta E$): Gaps near 2 eV provide strong charge-transfer and high β but risk red/near-IR absorption; gaps above 4 eV ensure visible transparency but reduce β. Targeting 2–4 eV, we minimize: $f\Delta E = \max(0, 2 - \Delta E) + \max(0, \Delta E - 4)$.
- Energy per Atom ($Etotal/Natoms$): Proxies thermodynamic stability. Minimizing this encourages practical synthesizable NLO materials. Positive values indicate SCF failures or strained geometries, so we reject those molecules.

Each algorithm evaluated ~2,000 molecules across 20 independent runs initialized from mutated common molecular scaffolds.

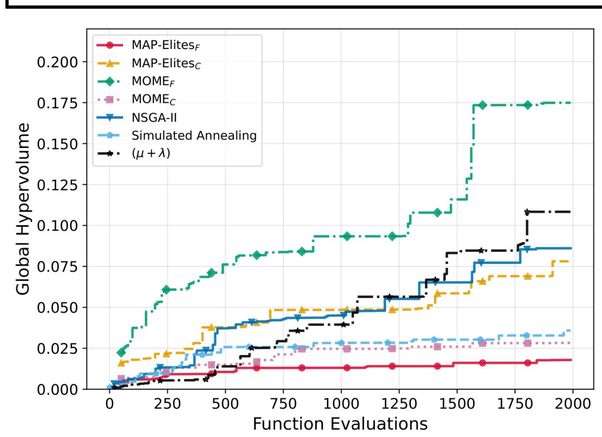

(a) Across Evaluations

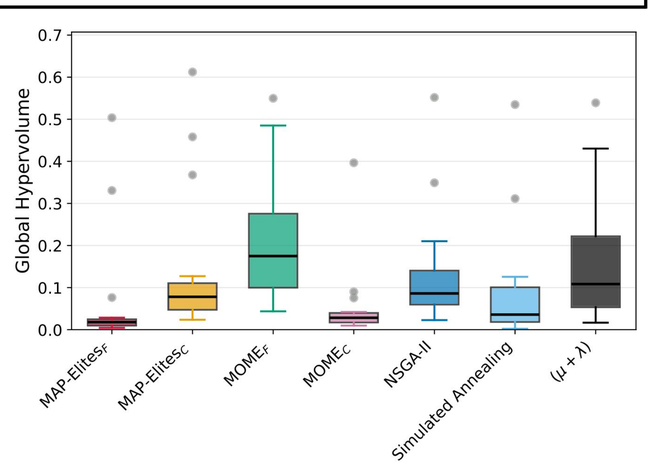

(b) Final Pareto Fronts

Figure 1: Global Hypervolume Scores Across 20 Runs of Each Algorithm: Objective scores were normalized based on the range across all 20 runs of all algorithms, but with chemically illegal/impossible scores discarded. (a) Median HV for each algorithm across function evaluations. MAP-Elites and MOME archives characterize molecular diversity by atom count (5 to 30) and bond count (4 to 32). The range of values is evenly divided along each dimension by type: C = coarse (10 × 10) or F = fine (20 × 20). MOME-F (MOME with the fine-grained binning) is the best. (b) Box-and-whisker plots of HV for final Pareto fronts. MOME-F and $(\mu+\lambda)$ both have high upper quartiles, and spread more across the range of higher scores.

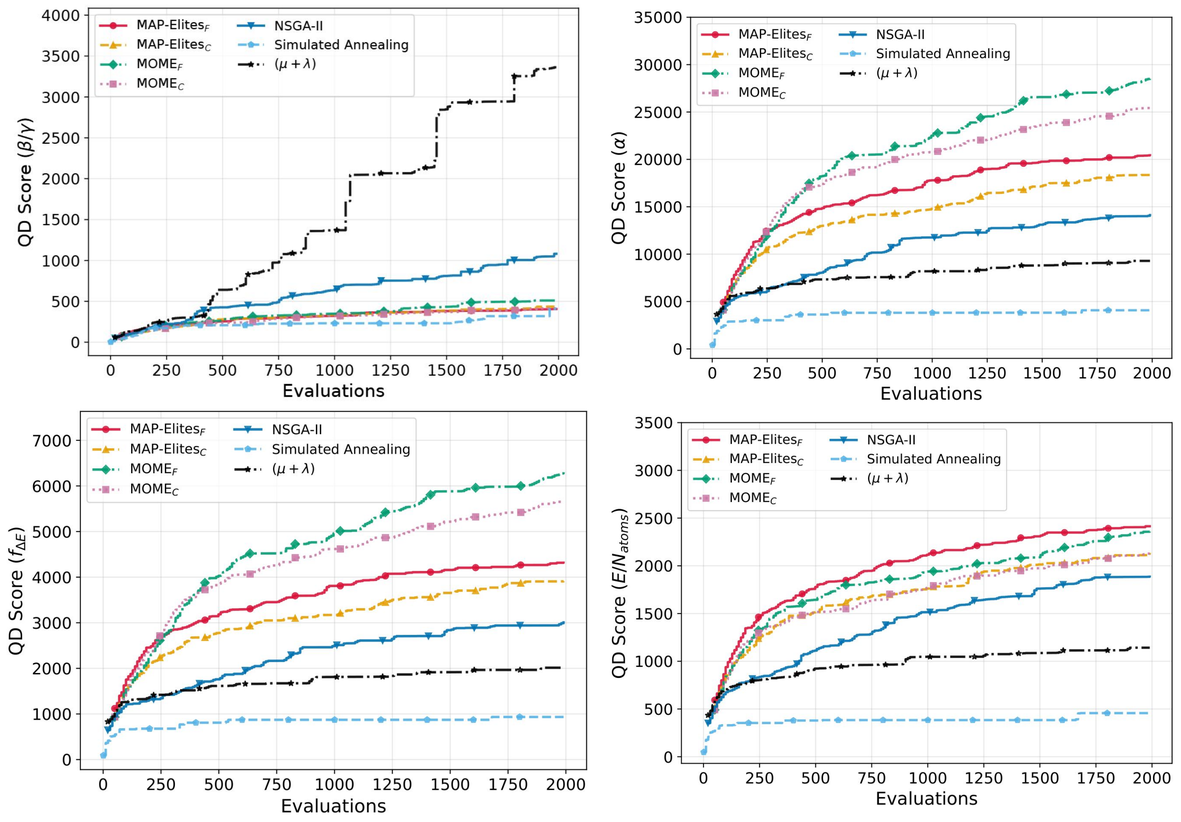

Figure 3: Fine-grained Archive Median QD Scores by Objective Across 20 Runs of Each Algorithm: QD scores are calculated with respect to exactly one objective in each figure. Median QD for β/γ is qualitatively similar to raw objective scores for β/γ. Median QD for $f\alpha$, both MOME approaches perform the best, with MAP-Elites approaches beneath them, and NSGA-II trailing close behind, beating $(\mu + \lambda)$ and simulated annealing. Median QD for $f\Delta E$ are qualitatively similar to the $(f\alpha)F$ results. Median QD for $Etotal/Natoms$, the MOQD and QD methods are more tightly clustered with NSGA-II near the top.

Figure 2: Median MOQD with fine-grained binning. Here MOME-F is clearly superior, followed distantly by MAP-Elites-C, then NSGA-II. All figures in this poster were generated using fine-grained binning. However, the results are qualitatively similar under coarse-grained binning and further demonstrate that evolution with fine-grained bins outperforms evolution with coarse-grained bins, even when evaluated using coarse-grained performance metrics. Coarse figures are included in the paper itself and along with all references are attached to the QR code at end of this poster.

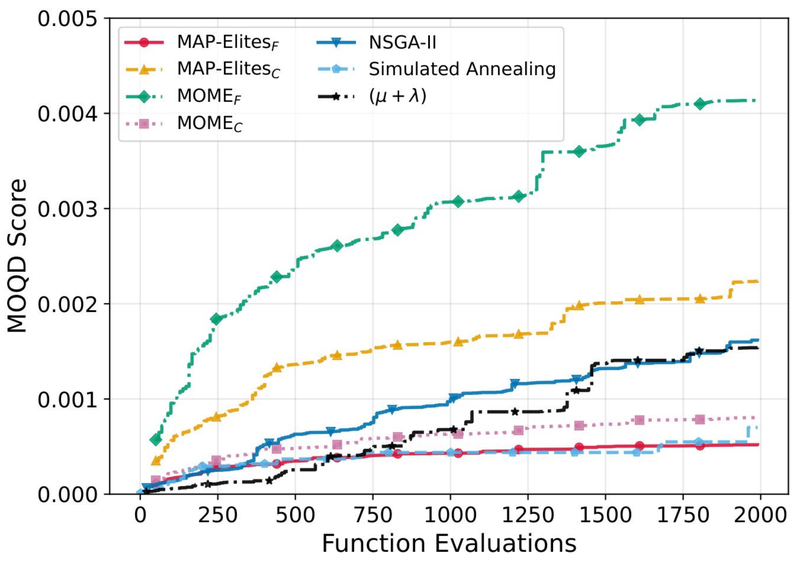

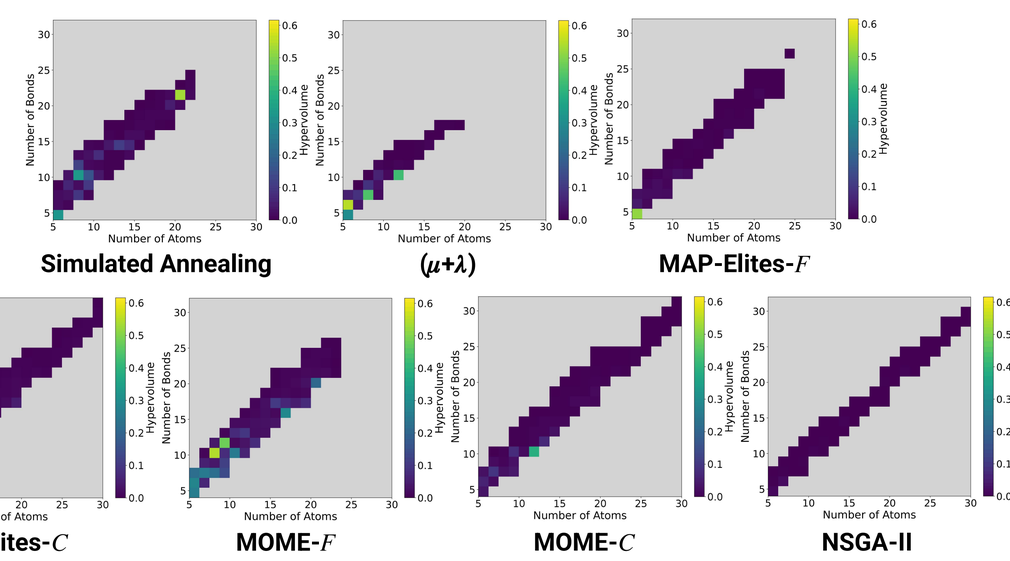

Figure 5: Fine-grained Mega Archive Hypervolume Heatmaps: Fine-grained archives that combine solutions from each algorithm across all 20 seeds, with heat scale showing each bin's HV score. The x-axis is the atom count and the y-axis is the bond count. Each portrays a diagonal slash through the archive, because bond count and atom count are correlated, but the length/width of the slash varies.

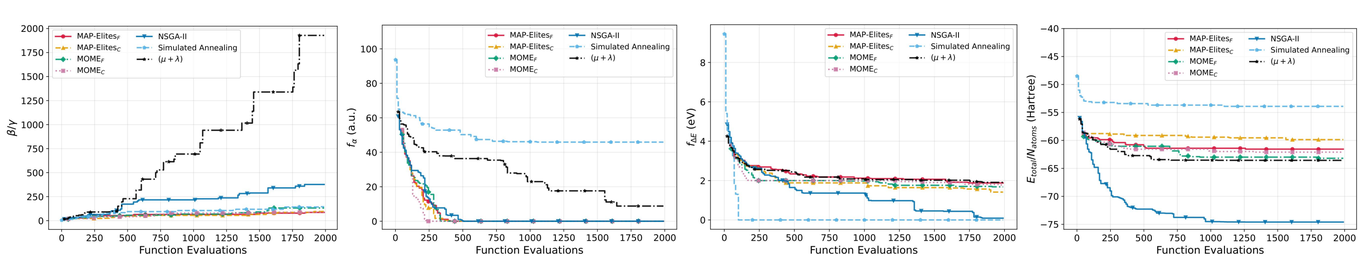

Figure 4: Median Best Objective Scores Across 20 Runs of Each Algorithm: Median β/γ, $(\mu + \lambda)$ outperforms all others, with NSGA-II second-best. Median $f\alpha$, all but simulated annealing and $(\mu + \lambda)$ reach a perfect score of 0. Median $f\Delta E$, NSGA-II and simulated annealing tie with perfect scores of 0. Median $Etotal/Natoms$, NSGA-II is the best at minimizing this objective, with most other methods clustering closer together, including the single-objective methods that were unaware of this objective. Only simulated annealing is exceptionally poor.

Conclusions & References

Results show that with a fine-grained archive, MOME covers the broadest range of structurally diverse molecules, and these molecules produce the largest global hypervolume. NSGA-II also produced decent hypervolume scores and performed well on individual objectives, but scores poorly on QD/MOQD metrics. The performance of $(\mu + \lambda)$ seemed good due to high (β/γ) scores, but this accomplishment came at the expense of the other objectives. These molecules are not useful NLO molecules. We will further investigate how to use these algorithms to guide the search for molecules with desired properties.

We extended this work by replacing arbitrary atom/bond-based descriptors with CVT archives; the arXiv paper is available via the QR code.

All code used to run our experiments is available at this url: https://github.com/DominicMashak/Molecular-Evolution

References and additional material