A Global-Novelty Reward for Multi-swarm Safety Checking: Steering Diversification at the Objective Level

Follow-up study #2 to Kumazawa et al., PAAMS 2023

2026

Abstract

The objective of MSPSO-SAFE rewards *within-path* diversity through the term *pathlen* · *firstvisit*, but nothing in the objective rewards visiting abstract states that are new *globally*, across the whole incremental search. We add a global-novelty bonus that pays a particle for each abstract state it reaches for the first time across all outer iterations. The bonus gives a small but consistent coverage improvement (+0.0021 in coverage AUC overall, positive for 4 of 5 process counts) that grows with its weight, and unlike the paper's counterexample-driven updates it acts directly on the coverage objective. We position novelty as the cheapest of three coverage levers (objective reward, restart point, metric), effective on its own and composable with the others. Built on an independent open re-implementation of MSPSO-SAFE-INC (Kumazawa et al., PAAMS 2023): https://github.com/yaskodama/kuma2023-repro.

1 Motivation

In Formula (3) of [1], diversification is encouraged only *inside* one decoded path: *firstvisit*(x) counts distinct abstract states on the path. Across outer iterations there is no term that distinguishes "this path discovered a brand-new abstract state" from "this path re-walked known states". Because MSPSO-SAFE-INC is a re-initialisation method, successive runs therefore tend to rediscover the same easily reachable abstract states, and coverage growth slows. Cheng et al. [2] show that explicit novelty/diversity promotion helps PSO escape such basins; we bring that idea to the model-checking objective.

2 Method: global-novelty bonus

Let V be the set of abstract states visited so far across the whole incremental run (maintained online). For a decoded path we add to Formula (3)

$$NOV(x) \;=\; \beta \cdot \big|\{\, a \in \mathrm{abs}(\mathrm{path}(x)) : a \notin V \,\}\big|,$$

the number of abstract states on the path that are globally new. The term is on the same scale as *pathlen* · *firstvisit* ($\leq$ *maxlen*), so with β of order one it can actually influence particle selection — in contrast to APUP, which only modifies a term bounded by 1 and is therefore inert regardless of magnitude.

3 Setup

Reproduced benchmark, n=1000 states, faithful Table-1 PSO budget (particle 20, generations 30, swarm cap 6). Coverage is the diversity-AUC (mean cumulative AC_{div} over 120 outer iterations), averaged over $p \in \{2, \ldots, 6\}$ and 3 RNG repetitions.

4 Results

The novelty bonus improves coverage by +0.0021 overall and is positive for four of five process counts (Table 1). The effect grows with the weight β (β=1 : +0.0014, β=3 : +0.0030 on a four-instance subset), confirming the term is the active cause rather than noise. The single regression (p=3) is an instance whose reachable abstract space is small, where extra diversification pressure buys little. Figure 1 shows the averaged p=5 curve.

Table 1: Diversity-AUC, novelty bonus (β=3) vs baseline.

p	2	3	4	5	6	overall
none	0.7329	0.5941	0.7040	0.8450	0.8113	0.7375
NOV	0.7357	0.5910	0.7086	0.8479	0.8144	0.7395
Δ	+.003	−.003	+.005	+.003	+.003	+.0021

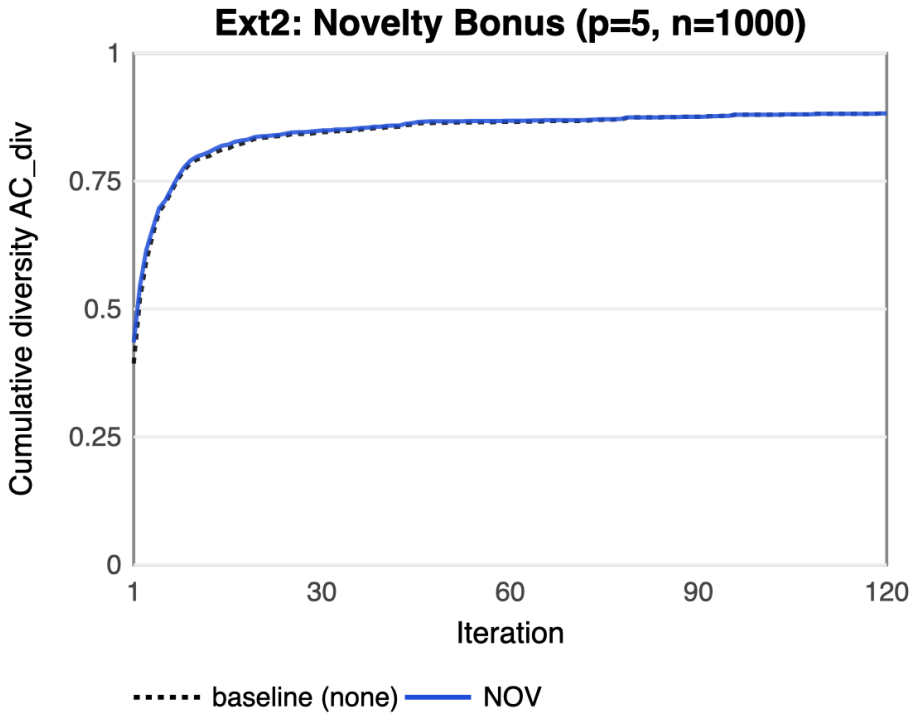

Figure 1: Averaged cumulative AC_{div} (p=5): the global-novelty bonus slightly accelerates coverage over the baseline.

5 Discussion

The novelty bonus is the most surgical of the coverage levers: it changes only the objective, needs no extra state beyond the visited-abstract set, and is the correct counterpart to the paper's APUP/APP — it rewards the right quantity (global novelty) at the right scale. Its effect is modest because, at the full Table-1 budget, a single run already covers much of the reachable abstract space; its value rises at smaller per-run budgets (cf. follow-up #4) and when composed with persistent restarts (follow-up #3).

6 Conclusion

A global-novelty term turns MSPSO-SAFE's within-path diversity reward into a whole-run one, giving a small consistent coverage gain that scales with its weight. It is cheap, composable, and conceptually the fix that APUP/APP should have been. Future work: adaptive β and behavioural (path-shape) novelty.

References

[1] T. Kumazawa, M. Takimoto, Y. Kodama, Y. Kambayashi. *Enhancing Safety Checking Coverage with Multi-swarm Particle Swarm Optimization.* PAAMS 2023.

[2] S. Cheng, Y. Shi, Q. Qin. *Promoting Diversity in Particle Swarm Optimization to Solve Multimodal Problems.* ICONIP 2011.