Peer Review Report: "Bio-inspired and Evolutionary Computing: Aspects of Project and Theory" by Richard Murdoch Montgomery

Manuscript ID / Type: Article (tutorial/review style with original implementations); submitted to a mathematical/computational periodical.
Reviewer: Grok (on behalf of the review team)
Date: 29 March 2026
Recommendation: Minor Revision – Accept pending targeted improvements. This is a strong, well-structured tutorial-style manuscript that would make an excellent educational or survey article in a journal that publishes computational intelligence reviews, algorithm tutorials, or "practical theory" pieces (e.g., *IEEE Transactions on Evolutionary Computation*, *Swarm Intelligence*, *Journal of Computational Science*, or similar).

1. Overall Assessment

The manuscript delivers a clear, self-contained introduction to core bio-inspired and evolutionary computing paradigms (Genetic Algorithms, Particle Swarm Optimisation, Ant Colony Optimisation, with mentions of neuro-evolution and multi-objective optimisation). It successfully bridges **theory** (mathematical formalisms), **practice** (Python implementations), **experimentation** (benchmark results with visualisations), and **critique** (discussion of strengths/limitations).

The level is appropriate for advanced undergraduates, graduate students, or practitioners new to the field. The writing is accessible yet sufficiently rigorous for a mathematical/computational audience. Original contributions include the provided working code, multi-panel convergence visualisations, and a side-by-side PSO vs. random-search comparison—valuable for teaching and reproducibility.

Strengths

- Logical flow: Abstract → Introduction → Methodology (math) → Results (experiments + figures) → Discussion → Conclusion + code appendices.
- Excellent balance of equations (correct and well-presented), pseudocode-level explanations, and full Python implementations.
- High pedagogical value: convergence plots, fitness landscapes, diversity metrics, and parameter discussions are textbook-quality.
- Reproducibility is high; the GA and PSO classes are complete and appear functional.
- References are current (up to 2024) and appropriately cited.

Major Concerns (must be addressed)

1. **Title typo / awkward phrasing**
   "Aspects of Project and Theory" is almost certainly intended to be "Aspects of **Practice** and Theory". Change this for professionalism.
2. **Incomplete coverage of ACO**
   Section 2.3 gives a full mathematical treatment of Ant Colony Optimisation, yet the Results and Attachments sections contain **zero** ACO experiments, plots, or code. Either:
   - Add a brief ACO implementation and results (e.g., TSP on a small graph), or
   - Explicitly state that ACO is presented for completeness but not implemented here, and move the section to an appendix or shorten it.
     As written, the methodology promises three algorithms but delivers results for only two.
3. **Numerical inconsistencies in PSO results**
   - Text (p. 9): "final best fitness of approximately 1.5×10^{-12}" on the sphere function.
   - Text (p. 10) & Table 1: 4.82×10^{-5}.
   - Random-search sphere value also mismatches (text says 1.26×10^{-1}; table says 1.26×10^{1}).
     These are likely different runs or slight code variations, but they must be reconciled or explained (e.g., "representative run with 30 particles, 100 iterations").
4. **GA limited to binary OneMax**
   Table 1 lists "GA_Real" as N/A for continuous benchmarks. The manuscript would be stronger if a real-valued GA variant (or DE) were added for fair comparison on sphere/Rastrigin/etc.

Minor Concerns (should be fixed)

- **Formatting / OCR artefacts** in the supplied PDF (e.g., "for mediaimage", "float(\nF)", missing closing quotes in strings). The final LaTeX/PDF version must be clean.
- **Table 1** caption and layout: "GA_Real" column is confusing when no real-valued GA is presented. Rename or remove.
- **References**: The Medium blog (Zirpe, 2024) is acceptable for accessibility but ideally replaced by a peer-reviewed source. All DOIs/URLs should be verified.
- **Figure captions**: Excellent, but ensure every multi-panel figure (Figs. 1–4) has sub-figure labels (a)–(f) explicitly matched in the caption.
- **Code style**: Minor transcription issues in the provided excerpts (e.g., `roulette_whee1` typo, `float(\nF)`). In the published version, provide clean, syntax-highlighted code (or GitHub link) and confirm it runs with current NumPy/Matplotlib.
- **Neuro-evolution and MOEA mentions** in the Introduction are promising but never revisited. Either expand with a short example or remove the forward-looking sentences to avoid raising unfulfilled expectations.

2. Scientific / Technical Rigor

- **Mathematics**: Equations for GA (selection, crossover, mutation), PSO (velocity/position update, inertia scheduling), and ACO (transition probability, pheromone update) are standard, correctly derived, and accurately typeset in KaTeX.
- **Benchmarks**: OneMax (GA) and sphere/Rastrigin/Rosenbrock/Ackley (PSO) are canonical choices. Fitness-landscape visualisations are particularly effective.
- **Experimental design**: Population sizes, iteration counts, and parameter values are reasonable for illustration. Diversity and improvement-rate plots add depth beyond simple convergence curves.
- **Limitations section**: Balanced and honest—covers lack of optimality guarantees, parameter sensitivity, and computational cost. Good citation support.

No factual errors in the algorithmic descriptions were found.