POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Latent Space Exploration for Racing Track Design using Novelty Search

Laurea Magistrale in Computer Science and Engineering - Ingegneria Informatica

Author: Milo Brontesi

Advisor: Prof. Daniele Loiacono

Co-advisor: Prof. Pierluca Lanzi

Academic year: 2025-2026

---

1. Introduction

A core challenge in game development is the resource-intensive nature of design and content production. Depending on the context, creating new content can become a major bottleneck in development cycles. Modern games demand varied, replayable content without incurring prohibitive production costs and time constraints. Procedural Content Generation (PCG) addresses this by automatically creating digital assets through algorithmic methods [5, 13]. This automation acts not merely as a substitute for human effort, but allows developers to rapidly prototype, test, and design at a scale beyond traditional production limits. Furthermore, the ability of PCG methods to algorithmically generate high-quality content enables the exploration of novel game mechanics and innovative design paradigms.

A key challenge in PCG is generating content that is both engaging and meaningfully varied. Traditional Search-Based PCG and standard optimisation algorithms typically aim to converge on a single, globally optimal solution. However, in game and level design, discovering a single global optimum is rarely the goal. To accommodate diverse player preferences and varying design needs, developers require a broad portfolio of high-quality solutions. Quality-Diversity (QD) optimisation techniques are specifically designed to address this requirement [3]. Rather than searching for a singular optimum, QD algorithms apply evolutionary techniques to explore the design space by discovering a vast, structured set of solutions that are simultaneously high-performing (quality) and fundamentally distinct from one another across a defined behavioural space (diversity).

This thesis investigates the application of Quality-Diversity algorithms to the **procedural generation of racing tracks**. Specifically, this research proposes a data-driven pipeline based on Novelty Search with Local Competition (NSLC). Instead of relying on hand-crafted heuristics, the system operates within a behavioural space that is autonomously learned and dynamically maintained through unsupervised representation learning, utilising the deep neural architecture of a Variational Autoencoder (VAE) [10].

2. Methodology

To autonomously discover a diverse set of racing circuits, the system integrates a geometry generator, a physics-based simulation environment, and an evolutionary pipeline that adapts its own definition of "diversity" over time.

2.1. Track Generation and Simulation

Our procedural approach synthesises tracks using Voronoi diagrams [6] and Catmull-Rom spline interpolation [1]. Generating boundaries from subsets of Voronoi cells ensures that the resulting tracks are structurally contiguous and form closed loops. The generated points are converted to XML formats natively supported by The Open Racing Car Simulator (TORCS) [14]. Mathematical closure solvers and engine patches are applied to eliminate discontinuities and position gaps inherent to TORCS's sequential generation logic.

Tracks are evaluated inside headless Dockerised [4] instances of TORCS. For each track, the system executes automated multi-agent race simulations. The fitness objective is established as the *Total Overtakes* recorded during the race, optimising the geometries to favor dynamic, arcade-style gameplay.

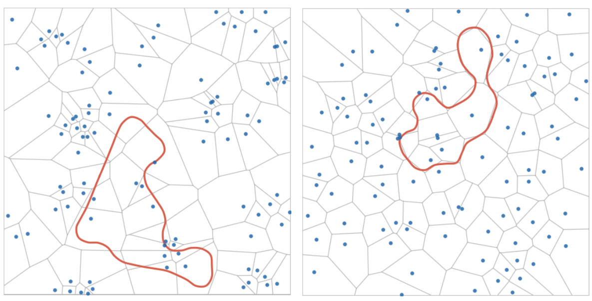

Figure 1: Two tracks generated starting from the Voronoi diagram.

2.2. Dynamic Representation Learning

Traditional QD algorithms require human experts to define the Behavioural Descriptors (BDs) manually. Inspired by the AURORA framework for autonomous skill discovery in robotics [2], our architecture delegates the task of mapping behavioural space to a Variational Autoencoder (VAE) [10].

The BD is derived not from static geometric shapes, but from the raw driving telemetry (speed, steering angle, and distance from the track boundary) extracted during a benchmark lap. To accommodate variable-length telemetry without zero-padding, the VAE leverages 1D Circular Convolutional Residual Blocks [9], preserving the continuous-loop topology of the track. A summary of the architecture is provided in Figure 2.

A major challenge in cyclic sequences is starting point bias. To guarantee a strictly shift-invariant embedding, the encoder projects the feature map via a custom Discrete Fourier Transform (DFT) Power Pooling layer, purposefully discarding phase information. This allows the VAE to map identical tracks with different starting coordinates to the exact same point in a dense, 32-dimensional latent space. The model is trained using a weighted cross-correlation aligned reconstruction loss coupled with cyclical KL divergence regularisation [7, 11]:

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \beta\,\mathcal{L}_{\text{KL}} \tag{1}$$

where the shift-aligned reconstruction error is defined as:

$$\mathcal{L}_{\text{rec}} = \|\mathbf{x}\|_{\mathbf{w}}^2 + \|\hat{\mathbf{x}}\|_{\mathbf{w}}^2 - 2\max_{\tau} \text{CC}(\tau) \tag{2}$$

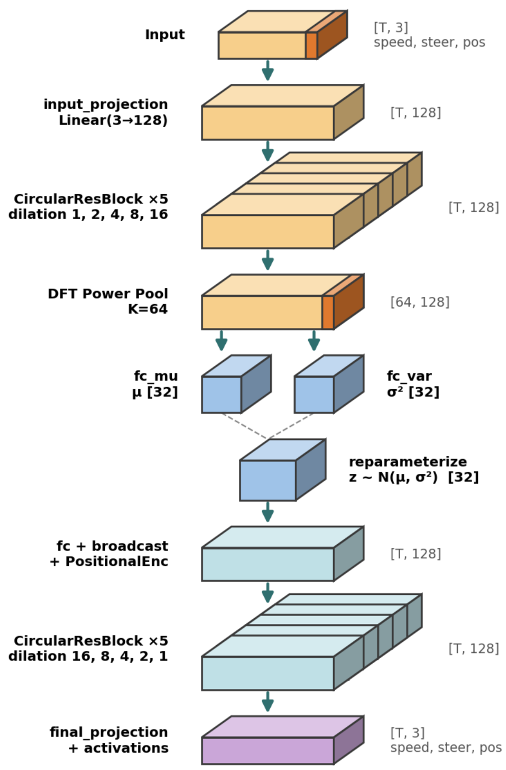

Figure 2: VAE architecture summary.

2.3. The Evolutionary Loop

The algorithm manages a continuous "Ask-Evaluate-Tell" sequence powered by the Pyribs framework [12]. Candidate geometries are evolved via structural mutation and spatial crossover. Once a track's telemetry is evaluated and its coordinates in the latent space are mapped by the VAE, it attempts insertion into an unstructured Proximity Archive. A candidate survives if its latent Euclidean distance to its k-Nearest Neighbours exceeds an adaptively scaled novelty threshold, or if it outcompetes local neighbours in overtakes [8].

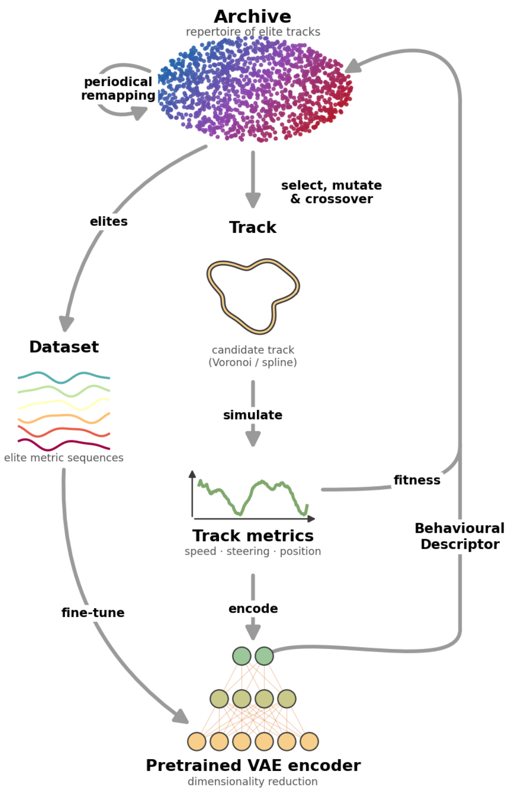

Figure 3: Evolutionary pipeline loop.

Periodically, a *Representation Update Routine* freezes the evolutionary search to fine-tune the VAE on the newly discovered elites. After fine-tuning, the archive is entirely re-encoded and remapped to the updated latent manifold. A summary of the pipeline is provided in Figure 3.

3. Experiments and Results

The VAE was initially pre-trained on a robust dataset of 20,000 diverse geometries, achieving high structural reconstruction quality with no posterior collapse across its 32 active dimensions.

To evaluate the pipeline, two primary 1,100-iteration configurations were executed and compared: `base-NS` (where the pre-trained VAE remains static) and `finetune-NS` (where the VAE is dynamically fine-tuned every 100 iterations).

3.1. Archive Growth and Fitness Improvement

Both setups utilised an adaptive novelty threshold targeting an archive size of roughly 500 elite solutions. The `finetune-NS` configuration exhibited a highly dynamic, saw-tooth archive growth curve, periodically purging sub-optimal solutions during latent space remapping cycles. Crucially, the dynamic configuration yielded a substantially faster acceleration in global fitness (as observable in Fig. 4). Refining the latent space actively consolidates high-performing elites, ensuring subsequent genetic operations breed from higher-quality populations.

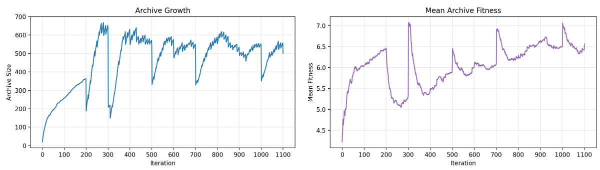

Figure 4: Archive growth and mean fitness graphs for a run of the `finetune-NS` configuration.

3.2. Mitigation of Clone Solutions

A major limitation observed in the `base-NS` configuration was the accumulation of geometric "clones". When the evolutionary search pushes deep into unseen design space, static VAEs often fail to properly cluster these novel solutions due to large unmapped voids in the pre-trained space. Consequently, minor genetic mutations generate highly similar tracks that the static VAE perceives as highly distinct, bypassing local competition and polluting the archive with redundancies.

In stark contrast, the `finetune-NS` architecture substantially mitigated this problem. By continually adapting the encoder to the newly discovered out-of-distribution trajectories, previously isolated "outlier" tracks are properly contextualised and clustered. This guarantees that only genuinely distinct geometries are rewarded with novelty scores, yielding an elite repertoire of drastically improved diversity.

3.3. Latent Manifold Stability

Despite the low volume of samples provided during periodic fine-tuning routines, the model managed to avoid catastrophic forgetting and posterior collapse. The KLD metric remained stable and strictly aligned with pre-training baselines. Visualisations of time-domain reconstructions confirmed that the dynamic VAE reliably captures novel, high-frequency behavioural phenomena that the static model failed to represent.

4. Conclusions

This research successfully bridges Quality-Diversity optimisation, unsupervised representation learning, and autonomous skill discovery to advance the procedural generation of functional game content. By defining diversity via physics-driven simulation telemetry rather than arbitrary mathematical shapes, the pipeline uncovers track designs inherently linked to experiential vehicle dynamics.
Deploying an AURORA-style architecture eliminated the dependency on human-crafted heuristics. Furthermore, comparative experiments indicated that dynamically fine-tuning the VAE is essential for sustained evolutionary search. Adapting the behavioural space curtails the survival of redundant track clones and ensures the algorithm's continuous mapping of the shifting phenotypic frontier.
Future work may aim to incorporate more advanced continual learning techniques into the VAE fine-tuning process to further stabilise the loss manifolds, and introduce more sophisticated multi-agent AI controllers to generate richer phenotypic data.

5. Source Code

This thesis is supported by a full implementation of the procedural content generation pipeline. The complete source code, including the Novelty Search implementation, track generation algorithms, Docker configurations, and telemetry analysis tools, are publicly available at the following GitHub repository:
`https://github.com/zibasPk/Novelty-Search-for-Racing-Track-Design`

References

[1] Edwin Catmull and Raphael Rom. A class of local interpolating splines. In ROBERT E. BARNHILL and RICHARD F. RIESENFELD, editors, *Computer Aided Geometric Design*, pages 317–326. Academic Press, 1974.

[2] Antoine Cully. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '19, page 81–89. ACM, July 2019.

[3] Antoine Cully and Yiannis Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259, 2018.

[4] Docker Inc. Docker: Accelerated container application development, 2026. Accessed: 2026-06-02.

[5] Mahdi Farrokhi Maleki and Richard Zhao. Procedural content generation in games: A survey with insights on emerging llm integration. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 20(1):167–178, November 2024.

[6] Steven Fortune. A sweepline algorithm for voronoi diagrams. *Algorithmica*, 2(1):153–174, 1987.

[7] Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 240–250, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

[8] Luca Grillotti and Antoine Cully. Unsupervised behaviour discovery with quality-diversity optimisation, 2022.

[9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[10] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[11] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.

[12] Bryon Tjanaka, Matthew C Fontaine, David H Lee, Yulun Zhang, Nivedit Reddy Balam, Nathaniel Dennler, Sujay S Garlanka, Nikitas Dimitri Klapsis, and Stefanos Nikolaidis. Pyribs: A bare-bones python library for quality diversity optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '23, page 220–229, New York, NY, USA, 2023. Association for Computing Machinery.

[13] Julian Togelius, Georgios N. Yannakakis, Kenneth O. Stanley, and Cameron Browne. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):172–186, 2011.

[14] Bernhard Wymann, Eric Espié, Christophe Guionneau, et al. TORCS: The open racing car simulator. `https://sourceforge.net/projects/torcs/`, 2026. Version 1.3.9, accessed 2026-05-28.