Queen Mary
University of London

The Search for Serendipity: novelty and illumination

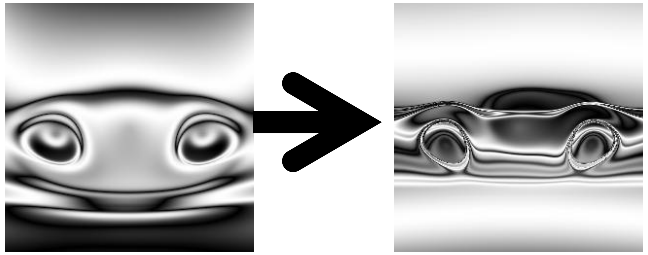
*

Laurissa Tokarchuk – laurissa.Tokarchuk@qmul.ac.uk

May 2018

Game AI Research Group
http://gameai.eecs.qmul.ac.uk

CIS centre for intelligent sensing

Cognitive Science Research Group
http://cogsci.eecs.qmul.ac.uk

*cluster006.ovh.net/~fabelier/wp-content/uploads/2012/05/stanley_nstalk.ppt

Novelty Search

(Lehman and Stanley 2008)

- An evolutionary algorithm which *abandons* objective fitness
- Objective fitness is replaced by promoting phenotypic *novelty*
- Fitness is assigned as the mean Euclidean distance between the solution and its K-Nearest Neighbours from the population and an archive of previous highly novel solutions

Example of a *deceptive* domain

Georgios Methenitis https://www.youtube.com/watch?v=JIQP15tt5AI

Novelty AND High Performance

Diversity in *Asteroids*

- Play Asteroids by linearly combining novelty search and objective fitness
- 5 different phenotypic descriptors
- 6 different ratios of NS to points scored

Results

- Pure NS performs sub-optimally
- Small proportion of NS has no effect on points scored

However: High performing individuals in NS populations show *higher levels of diversity* than objective fitness

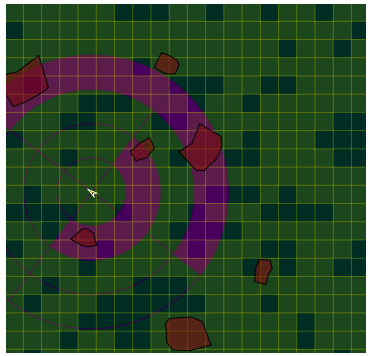

GC_μ - ground covered

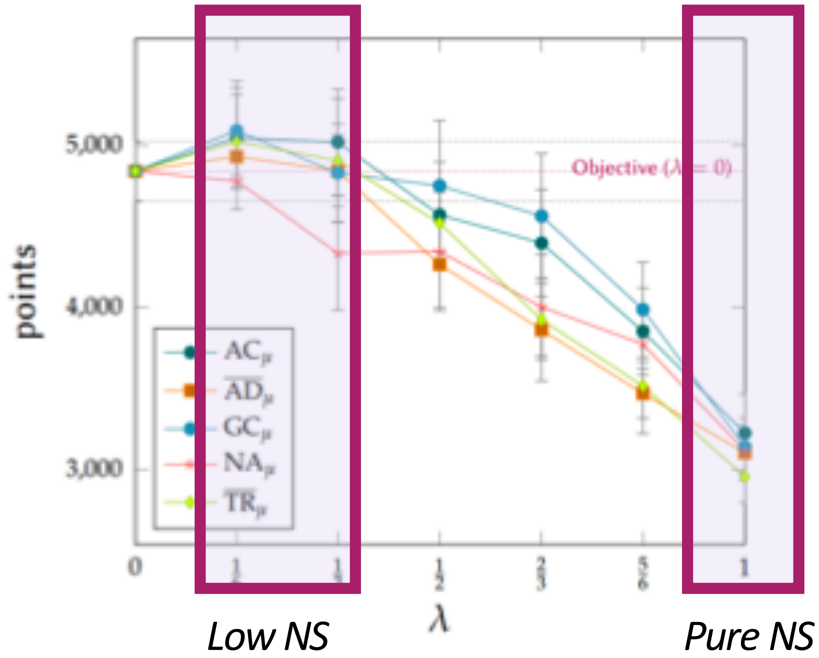

Smith, Tokarchuk and Fernando (2015)

Evolving diverse strategies through combined phenotypic novelty and objective function search

Multiple Objectives & Encouraging Diversity

Multiple Criteria Novelty Search (Lehman & Stanley 2010)

- MCNS: All individuals meet a defined minimal criteria or novelty = 0.

Multiple Assessment Directed Novelty Search
(Smith, Tokarchuk and Wiggins, 2016)

- MADNS: Highest performing individuals for each objective given the same fitness as the most novel individual.
- MCMADNS: MADNS + MCNS

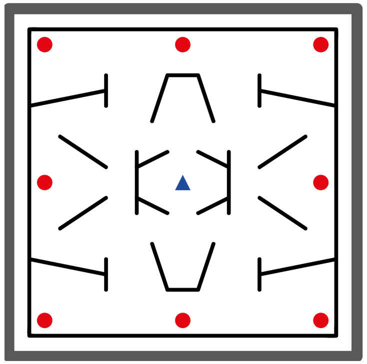

Results (*highlights*)

- MADNS outperforms NS as phenotypic landscape increases.
- MCNS helps to restrict exploration, MCMADNS assistswhile retaining some benefits of divergent exploration.

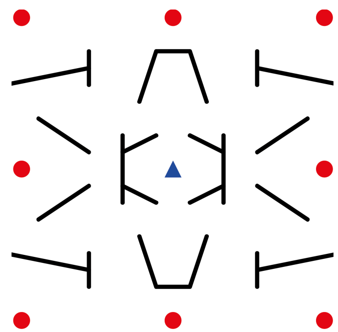

MADNS - Smith, Tokarchuk and Wiggins, 2016
Exploring Conflicting Objectives with MADNS: Multiple Assessment Directed Novelty Search (GECCO 16) &
Harnessing Phenotypic Diversity towards Multiple Independent Objectives (GECCO 16)

Illumination Algorithms

Motivation:

- Neither random sampling in big search spaces or finding fitness peaks by chance is effective in large state spaces.

Goal:

- Illuminate the fitness potential of each region of the feature space

Evolutionary algorithms which produce wide range of behavioural niches:

- optimised solution(s)
- high performing
- diverse

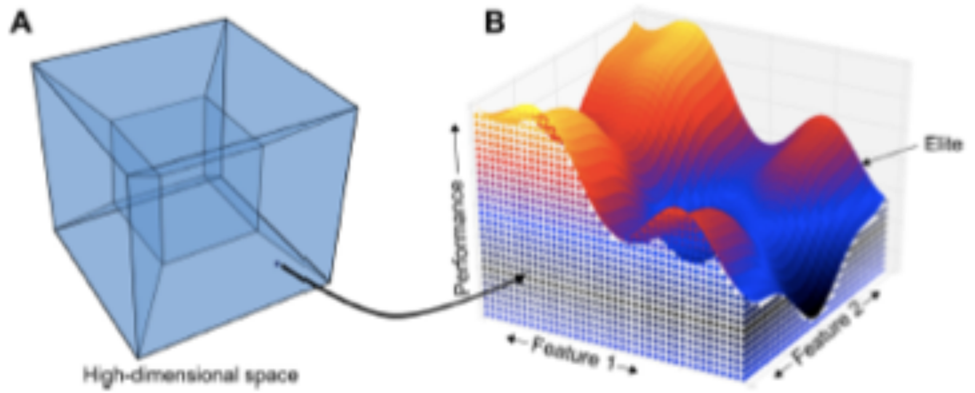

SHINE (Smith, Tokarchuk and Wiggins, 2016)

Rapid Phenotypic Landscape Exploration through Hierarchical Spatial Partitioning (PPSN 16)

MAP-Elites (Mouret & Clune, 2015)

Spatial Hierarchical Illuminated NeuroEvolution (SHINE)

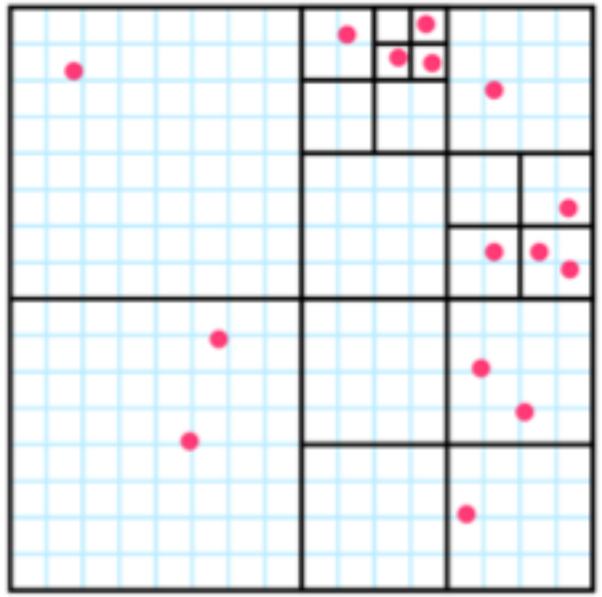

An illumination algorithm using hierarchical spatial partitioning

- Uses an n-dimensional tree
- Stores ***high performing individuals in each leaf***
- High performing individuals can be calculated in a different ways
- Number of individuals to store in each leaf vertex is determined by the depth within the tree
- Subsequent generations randomly assigned using proportional selection weighted by *1/depth within the tree*

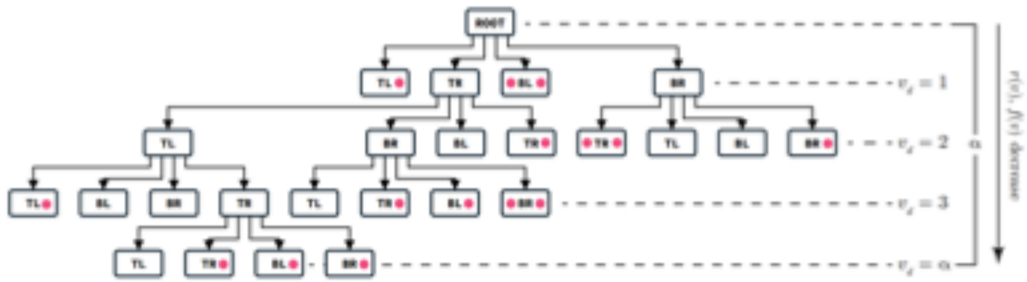

SHINE

Maze Navigation Experiments

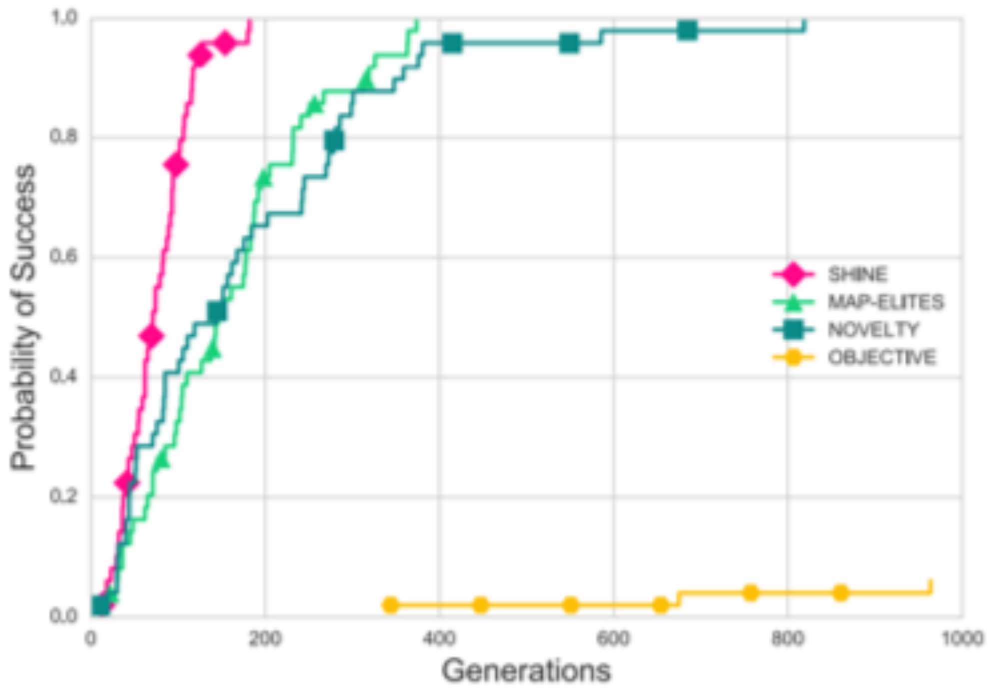

Hard maze (single exit)

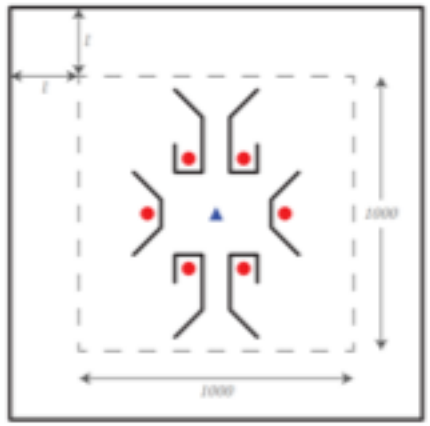

Large scaled maze
(multiple exit, l = 10000)

Results

- SHINE significantly outperforms all other algorithms (including MADNS) in objective fitness
- SHINE has higher level of domain coverage than MAP-Elites and NS
- Similar levels of exploration uniformity

Conclusions

- Promoting exploration of the phenotypic landscape is beneficial in overcoming deception and producing diverse sets of solutions to objectives
- Capable of developing more diverse and interesting controllers for applications such as video games
- As tasks in areas such as robotics become more complex, these methods may be beneficial in reducing fragility and increasing adaptability of controllers