

# Exploring Conflicting Objectives with MADNS: Multiple Assessment Directed Novelty Search

Davy Smith, Laurissa Tokarchuk and Geraint Wiggins

Queen Mary University of London

School of Electronic Engineering and Computer Science

London, E1 4NS, UK

{david.smith,laurissa.tokarchuk,geraint.wiggins}@qmul.ac.uk

## ABSTRACT

Novelty search is an evolutionary approach which promotes phenotypic diversity in a population. Novelty search has been successfully applied to a wide range of domains and a number of variants have been proposed. Here we introduce Multiple Assessment Directed Novelty Search (MADNS), which exploits the notion that a diverse population optimised through phenotypic novelty may contain solutions to multiple conflicting objectives. We show that by utilising the MADNS algorithm, an evolutionary trajectory may be simultaneously directed towards conflicting objectives. We conclude that, through applying MADNS and MC-MADNS, a divergent evolutionary trajectory may be directed to provide simultaneous solutions to multiple conflicting problems in domains with large potential for exploration.

## Keywords

Novelty search; algorithm design; phenotypic diversity; neuroevolution; evolutionary robotics

## 1. INTRODUCTION

Novelty search [?] is an evolutionary procedure which optimises a population through the promotion of behavioural diversity. Unlike traditional fitness based search, which converges towards a particular solution, novelty search is a divergent procedure, exploring the phenotypic landscape for potentially useful solutions in a given domain.

Here we propose a novel approach — multiple assessment directed novelty search (MADNS), which allows the trajectory of phenotypic exploration to be directed simultaneously towards multiple conflicting objectives. We propose variants of the MADNS algorithm based on both traditional novelty search and minimal criteria novelty search (MC-MADNS), highlighting the suitability to extend to unbounded domains. Our results show that, in large phenotypic landscapes, directing novelty search towards multiple conflicting objectives produces more optimal solutions and

with greater reliability than both novelty alone and minimal criteria novelty search (MCNS).

## 2. MULTIPLE ASSESSMENT DIRECTED NOVELTY SEARCH

MADNS is an extension to the novelty search algorithm that rewards high performing solutions over a number of predefined objectives. Formally, let the set  $P$  denote the current population, with an individual solution defined as  $\rho \in P$ . Next, for a given domain with  $k$  objectives, defined by the functions  $a_1(\rho), a_2(\rho), \dots, a_k(\rho)$ , where  $a_k : P \mapsto \mathbb{R}$ , let the set  $A = \{a_1(\rho), a_2(\rho), \dots, a_k(\rho)\}$ . Let the subset  $Q \subset P$  contain the maximal solutions for each objective, where:

$$Q = \arg \max_{\rho \in P} a_1(\rho) \cup \arg \max_{\rho \in P} a_2(\rho) \cup \dots \cup \arg \max_{\rho \in P} a_k(\rho).$$

Let  $f_{nov}(p)$  be the novelty of a solution and the maximal novelty value for the current population be defined as  $\alpha = \max_{\rho \in P} f_{nov}(\rho)$ . Finally, let  $f_{mad}(\rho)$  be the fitness of an individual solution, calculated as in ??:

$$f_{mad}(p) = \begin{cases} f_{nov}(p) & \text{if } p \notin Q, \\ \alpha & \text{if } p \in Q. \end{cases} \quad (1)$$

A minimal criteria variant of this (MC-MADNS) may be defined through the replacement of novelty search with MCNS [?]:

$$f_{mc-mad}(p) = \begin{cases} f_{mcns}(p) & \text{if } p \notin Q, \\ \alpha & \text{if } p \in Q. \end{cases} \quad (2)$$

## 3. EXPERIMENT

Our experimental domain is based upon previous studies which have assessed novelty search and variants of the algorithm [?]. The task domain is a simulated maze, in which an agent controller must navigate from an initial starting-point to one of a possible number of exit points within a fixed time limit. To fully assess both novelty search and MCNS, the tested domains include both bounded and unbounded versions (??).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '16 July 20-24, 2016, Denver, CO, USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4323-7/16/07.

DOI: <http://dx.doi.org/10.1145/2908961.2908975>




Figure 1: Mazes tested.

## 4. RESULTS

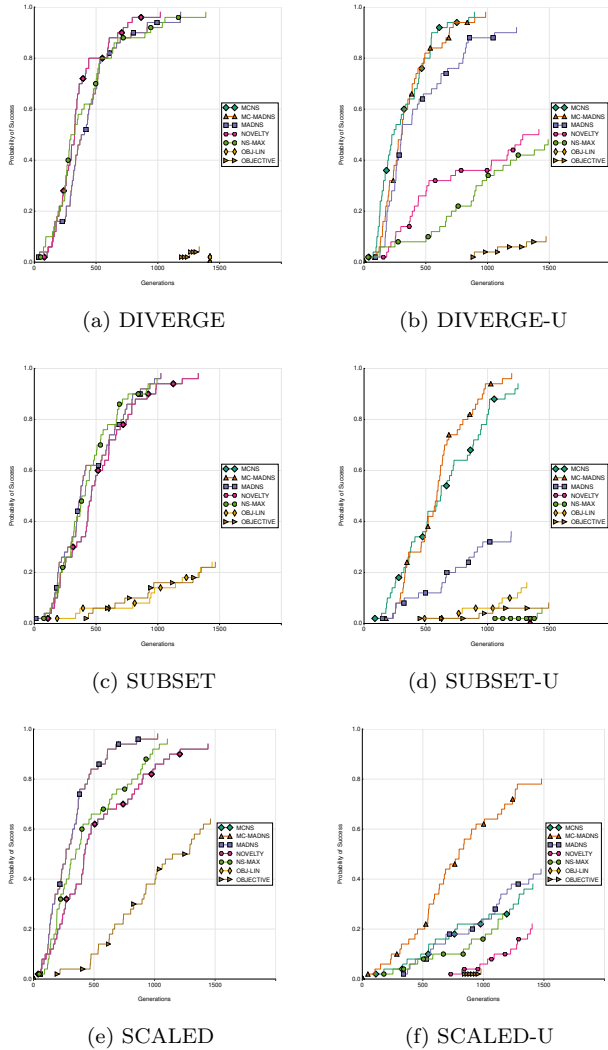


Figure 2: Probability of locating all exits within number of generations

?? shows the probability of success for each algorithm in each of the mazes over 100 evolutionary runs. In the bounded domains, the probability of success between MADNS and NS and their minimal criteria variants, MC-MADNS and MCNS, is identical due to the impossibility of a solu-

tion to fail the minimal criteria. In the DIVERGE domain, NS achieves its maximum value in fewer generations than MADNS. Conversely, in the SUBSET domain, MADNS achieves the same probability of success as NS in fewer generations. NS-MAX produces comparable levels of performance to NS and MADNS in all domains. In the SCALED domain, MADNS and MC-MADNS significantly outperform all other algorithms ( $p < 0.01$ ).

The probability of success achieved by MCNS and MC-MADNS in the DIVERGE-U domain is of a level comparable to the DIVERGE domain (????). In the SUBSET-U domain, MC-MADNS and MCNS achieve higher probability of success than in the SUBSET domain (????), with MC-MADNS outperforming MCNS in this domain (??). All algorithms produce results with lower probability of success (*pos*) in the SCALED-U domain MC-MADNS (*pos* = 0.80) and MADNS (*pos* = 0.43) are less effected by the unbounded SCALED-U domain than MCNS (*pos* = 0.38) and NS (*pos* = 0.21), with MC-MADNS significantly outperforming all other algorithms ( $p < 0.0001$ ).

## 5. CONCLUSION

In this paper, we presented a novel approach for hybrid directed-divergent search towards multiple conflicting objectives in both bounded and unbounded domains. Our results have shown that, in domains which have a low exploration potential, NS is an effective method for the optimisation of populations towards multiple conflicting objectives. However, our proposed MADNS and MC-MADNS, dependant upon the domain, are significantly more effective. For domains which have high exploration potential, the MADNS and MC-MADNS algorithms have a higher probability of success than both NS and MCNS.

### 5.1 Acknowledgements

This work was funded by EPSRC through the Media and Arts Technology Programme, an RCUK Doctoral Training Centre EP/G03723X/1.