Designing Fitness Functions for Odour Source Localisation

João Macedo
University of Coimbra
Coimbra, Portugal
jmacedo@isr.uc.pt

Lino Marques
University of Coimbra
Coimbra, Portugal
lino@isr.uc.pt

Ernesto Costa
University of Coimbra
Coimbra, Portugal
ernesto@dei.uc.pt

ABSTRACT

Locating odour sources is a hard task that has been addressed with a large variety of AI methods to produce search strategies with different levels of efficiency and robustness. However, it is still not clear how to evaluate those strategies. Simply evaluating the robot's ability to reach the goal may produce deceptive fitness values, favouring poor strategies that do not generalise. Conversely, including prior knowledge may bias the learning process. This work studies the impact of evaluation functions with various degrees of prior knowledge, in evolving search strategies. The baseline is set by performing multiple evaluations of each strategy with a function that only evaluates the task efficiency. A function was found that is able to produce strategies with equivalent performance to those of the baseline, whilst performing a single evaluation.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **Theory of computation** → **Design and analysis of algorithms**; • **Computer systems organization** → **Robotics**;

KEYWORDS

Fitness Evaluation, Evolutionary Robotics, Genetic Programming

ACM Reference Format:
João Macedo, Lino Marques, and Ernesto Costa. 2021. Designing Fitness Functions for Odour Source Localisation. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459524

1 INTRODUCTION

The natural phenomena associated with gas dispersion, make locating odour sources in real environments a hard task which is often not successfully solved by hand-designed strategies. Over the years, Artificial Intelligence (AI) and Evolutionary Robotics (ER) methods have been proposed to enable robots to learn and adapt search strategies to dynamic environments. However, most methods guide their learning process through fitness functions which are not easily designed for Odour Source Localisation (OSL). The stochastic nature of airflow and chemical dispersion lead to high degrees of uncertainty in the evaluation [2], being common for an individual to receive different fitness values over multiple evaluations. While an individual could be evaluated through the duration of the evaluation, such approach enables poor strategies to attain good fitness by chance. A common approach in ER is to evaluate each candidate solution multiple times to reduce the effects of uncertainty [2]. Unfortunately, the time associated with fitness evaluations render such approaches less desirable.

The present work focuses on accurately measuring the quality of an OSL strategy from a single evaluation. Ideally, the fitness function should evaluate how well a given strategy *searches* for the source, rather than evaluating how well it *finds* the source.

2 FITNESS FUNCTIONS FOR OSL

Most OSL works differ in the evaluation functions, but three aspects are usually present: (1) time spent; (2) distance travelled; (3) final distance to the source. Some works [1] reward the time spent sensing odour or the mean odour concentration sensed, which introduces a piece of prior knowledge: the concept that the agent should stay in the plume to track it. The knowledge that the wind carries the odour away from its source could also be introduced, rewarding upwind movements when sensing odour (u). Conversely, when the robot loses the plume, it either: (1) moved too much crosswind; (2) moved past the chemical source without finding it; or (3) the meandering of the plume moved it away from the robot. Drawing inspiration from biological strategies, the robot should be rewarded for moving crosswind or downwind to re-encounter the plume (l). The resulting fitness function is defined as:

$$F = \alpha \frac{d}{D} + \beta \frac{t}{T} + \gamma \left(1 - \frac{t_p}{t}\right) + \zeta(1-u) + \rho(1-l) \tag{1}$$

being d and D respectively the final and maximum distances to the source, t the time spent, T the maximum evaluation time, t_p the time spent sensing odour, α, β, γ, ζ and ρ weight coefficients and u and l computed as follows:

$$u = \sum_{i=1}^{N} \left(\frac{d_i \cdot \cos(uw_i)}{v \cdot \Delta t + 1} \right), l = \sum_{i=1}^{L} \left(\frac{d_i \cdot (\cos(cw_i) + \cos(dw_i))}{2v \cdot \Delta t + 1} \right) \tag{2}$$

being N and L respectively the amount of steps sensing and not sensing odour, v the robot's linear speed, Δt the duration of the control step and d, uw_i, cw_i, dw_i the distance travelled and the perceived upwind, crosswind and downwind directions in step i.

3 EXPERIMENTAL SETUP

3.1 Geometric Syntactic Genetic Programming

In this work, Geometric Syntactic Genetic Programming (GSynGP) [5] is used to evolve the robotic search strategies in the form of decision trees. The terminal and function sets are adapted from [5]. A population of 100 individuals is evolved for 100 generations. The trees are created with a maximum depth of 5. On each generation,

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07.
https://doi.org/10.1145/3449726.3459524

Table 1: Weights and objectives of the evaluation methods

Function	α	β	γ	ζ	ρ
F_1	1.0	1.0	0.0	0.0	0.0
F_2	1.0	1.0	1.0	0.0	0.0
F_3	1.0	1.0	0.0	1.0	1.0
F_4	1.0	1.0	1.0	1.0	1.0
LPP_1	d/D	t/T	$1-tp/T$	$1-u$	$1-l$
LPP_2	t/T	$1-u$	$1-l$	$1-tp/T$	d/D
LPP_3	$1-tp/T$	t/T	d/D	$1-l$	$1-u$
LPP_4	$1-l$	$1-u$	$1-tp/T$	t/T	d/D

10 random and 10 elitist immigrants are injected into the population. The parents are selected through binary tournaments, crossover is applied with 0.7 rate and with a number of iterations ranging from 1 to half of distance between the parents. Mutation is applied with 0.3 rate. The survivors selection is generational with an elite of 3.

3.2 Evaluation Environment

A simulation environment is devised based on simulator developed for ER and OSL [4]. It consists of a 40x40m arena with no obstacles and a single odour source. The environment is characterised by three variables: the wind speed (W_s, set to 1 m/s), the std. dev. of the Gaussian noise of the wind vectors (W_v, set to 0.1 rad) and the filament emission rate (F_r, set to 1 Hz). The initial position of the robot (x,y) is randomly selected from $N(36, 1.5)m$ in the odd-numbered evaluations. In the even-numbered evaluations, its y coordinate is sampled from $N(4, 1.5)m$. The position of the odour source is sampled uniformly, with $x \in [10, 12]m$ and $y \in [19.2.21.2]m$. A simulation is successful if the robot reaches a position less than 2m from the source or fails if the 1500s time limit is reached.

3.3 Evaluation Functions

The devised evaluation functions are presented on Table 1. Due to the difficulty in assigning proper weights to the components of the evaluation function, Lexicographic Parsimony Pressure [3] (LPP) is used to create four additional variants, differing in the order of the objectives. For completeness, a Pareto-based method (P) is also used to evolve the search strategies. The baseline for comparison (F_b) is set by evaluation each individual three times with F_1.

4 EXPERIMENTAL RESULTS

30 independent runs of GSynGP are performed for each function, being the resulting search strategies re-evaluated in the 30 instances of the simulation environment. The functions are compared through the Success Rate (S_r), the ratio of time spent in successful experiments (T_s), the ratio of final distance to the odour source in the unsuccessful evaluations (D_u), and the trajectory diversity (T_d), which is an indication of the bias of the evolutionary process. The higher the bias, the lower the trajectory diversity should be. The Wilcoxon test is applied to compare each function to F_b at a 95% confidence interval (the significance value is adjusted to 0.0056 with the Bonferroni correction). Table 2, shows that evaluating each individual multiple times (F_b) effectively reduces the uncertainty of the fitness value, raising S_r and lowering D_u. The Wilcoxon test showed that, F_b produces strategies significantly more robust than F_1 (p=0.004), F_3 (p=0.002), LPP_1 (p=0.003) and LPP_4 (p=0). F_2 is the only function with higher S_r than F_b (even though it is not significantly different (p=0.504)) and to produce a strategy with 100% S_r. The only strategies significantly faster (lower T_s) than those of F_b were produced by LPP_2 (p=0). Conversely, the strategies produced by LPP_3 are significantly slower than those of F_b (p=0). For D_u, only two significant differences were found: F_1 and LPP_2 are worse than F_b (p=0.002, 0.004). Finally, only P produced significantly less diverse trajectories than F_b (p=0). F_1 along with all LPP approaches produced significantly more diverse trajectories than F_b (all p-values=0).

Table 2: Performance of the evolved strategies

	Final strategies				Best strategy			
Function	S_r	T_s	D_u	T_d	#	S_r	T_s	D_u
F_b	52.7%	32.4%	48.3%	88.1%	2	90%	30.6%	40.5%
F_1	32.3%	24.3%	55.6%	90.1%	1	70%	59.8%	42.6%
F_2	55.3%	37.5%	47.0%	87.0%	1	100%	34.7%	n/a
F_3	30.8%	33.8%	52.4%	88.1%	1	93%	21.7%	39.0%
F_4	41.8%	32.4%	51.1%	88.5%	2	97%	28.2%	29.8%
LPP_1	33.1%	26.3%	53.8%	95.7%	1	87%	25.2%	24.2%
LPP_2	34.7%	22.0%	54.0%	96%	1	70%	22.0%	38.0%
LPP_3	33.9%	46.5%	44.1%	95.4%	1	87%	37.5%	60.4%
LPP_4	24.2%	65.3%	52.0%	93.3%	2	77%	80.7%	39.6%
P	44.8%	29.9%	53.4%	84.8%	1	97%	36.6%	46.9%

5 CONCLUSIONS

This paper designed and compared different evaluation functions for OSL strategies. The baseline was set by performing three evaluations of each individual with a function that measures the final distance to the source and the time spent to reach it. The results showed that the simplest function with prior knowledge (F_2) produces search strategies that are equivalent to those of F_b under all performance metrics, whilst making a single evaluation and consequently cutting down the computational time to one third. Moreover, F_2 produced a strategy more robust than the best of F_b (100% vs 90% success rate) and with equivalent speed.

In the future, methods to compute the weight coefficients should be studied and other concepts of prior knowledge should be explored and included into the evaluation functions.

ACKNOWLEDGEMENTS

This work was partially supported by the Portuguese Foundation for Science and Technology (FCT), projects UID/EEA/00048/2020, UID/CEC/00326/2020 and Ph.D. studentship SFRH/BD/129673/2017, co-funded by the European Social Fund, through the Regional Operational Program Centro 2020.

REFERENCES

[1] G.C.H.E. de Croon, L.M. O'Connor, C. Nicol, and D. Izzo. 2013. Evolutionary robotics approach to odor source localization. *Neurocomputing* 121 (2013), 481 – 497. Advances in Artificial Neural Networks and Machine Learning.

[2] Y. Jin and J. Branke. 2005. Evolutionary optimization in uncertain environments-a survey. *IEEE Transactions on evolutionary computation* 9, 3 (2005), 303–317.

[3] S. Luke and L. Panait. 2002. Lexicographic parsimony pressure. In *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*. 829–836.

[4] J. Macedo, L. Marques, and E. Costa. 2019. A Comparative Study of Bio-Inspired Odour Source Localisation Strategies from the State-Action Perspective. *Sensors* 19, 10 (2019), 2231.

[5] J. Macedo, L. Marques, and E. Costa. 2020. Locating Odour Sources with Geometric Syntactic Genetic Programming. In *European Conference on the Applications of Evolutionary Computation*. Springer.