



Search-based Generation of Waypoints for Triggering Self-Adaptations in Maritime Autonomous Vessels

Karoline Nylænder
karolinen@simula.no

Simula Research Laboratory and
Oslo Metropolitan University
Oslo, Norway

Shaukat Ali

shaukat@simula.no
Simula Research Laboratory
Oslo, Norway

Aitor Arrieta

aarrieta@mondragon.edu
Mondragon University
Mondragon, Spain

Paolo Arcaini

arcaini@nii.ac.jp
National Institute of Informatics
Tokyo, Japan

ABSTRACT

Self-adaptation in maritime autonomous vessels (AVs) enables them to adapt their behaviors to address unexpected situations while maintaining dependability requirements. During the design of such AVs, it is crucial to understand and identify the settings that should trigger adaptations, enabling validation of their implementation. To this end, we focus on the navigation software of AVs, which must adapt their behavior during operation through adaptations. AVs often rely on predefined waypoints to guide them along designated routes, ensuring safe navigation. We propose a multi-objective search-based approach, called WPgen, to generate minor modifications to the predefined set of waypoints, keeping them as close as possible to the original waypoints, while causing the AV to navigate inappropriately when navigating with the generated waypoints. WPgen uses NSGA-II as the multi-objective search algorithm with three seeding strategies for its initial population, resulting in three variations of WPgen. We evaluated these variations on three AVs (one overwater tanker and two underwater). We compared the three variations of WPgen with Random Search as the baseline and with each other. Experimental results showed that the effectiveness of these variations varied depending on the AV. Based on the results, we present the research and practical implications of WPgen.

CCS CONCEPTS

• **Software and its engineering** → **Search-based software engineering.**

KEYWORDS

autonomous vessels, simulation-based testing, multi-objective search, software testing

ACM Reference Format:

Karoline Nylænder, Aitor Arrieta, Shaukat Ali, and Paolo Arcaini. 2025. Search-based Generation of Waypoints for Triggering Self-Adaptations in Maritime Autonomous Vessels. In *Genetic and Evolutionary Computation*



This work is licensed under a Creative Commons Attribution 4.0 International License.
GECCO '25, July 14–18, 2025, Málaga, Spain
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1465-8/2025/07
<https://doi.org/10.1145/3712256.3726399>

Conference (GECCO '25), July 14–18, 2025, Malaga, Spain. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3712256.3726399>

1 INTRODUCTION

Self-adaptation enables systems to handle unforeseen situations while maintaining dependable operation. Implementing such self-adaptation in robotic software is a key objective of a large European project (RoboSAPIENS), within the framework of which the work reported in this paper is being conducted. The project focuses on extending the well-established MAPE-K self-adaptation loop [4] and applying it to various real-world industrial robotic software systems, enabling them to adapt effectively to unexpected situations. This paper focuses on one real-world application provided by the project, i.e., maritime autonomous vessels (AVs).

This paper considers the *navigation software* of the AV, as this is one of the most critical of its software components. This software typically relies on a navigation method using waypoints to help the AV accurately and safely generate its path to the intended destination [13]. When an AV follows its generated path, various conditions, such as environmental factors, sensor errors, or software malfunctions, can require it to adapt its behavior. Another type of event that may require adaptation is the presence of waypoints different from those the AV was supposed to follow. Indeed, the navigation software may not be properly tuned to handle these waypoints, and it may require an adaptation. This paper explicitly focuses on identifying waypoints that should trigger adaptations to prevent navigation software from malfunctioning. In particular, we propose a multi-objective search approach, called WPgen, to generate a set of waypoints by introducing minor modifications to an existing set of waypoints that should trigger adaptations. To guide the search, we defined two objectives, i.e., minimizing the variations of the original waypoints and maximizing the path length to generate waypoints. We employed NSGA-II as the multi-objective search algorithm in WPgen and implemented three strategies to seed its initial population, resulting in three variations of WPgen.

Our work focuses on generating minor variations to predefined waypoints and is thus related to test scenario generation for AVs. To this end, scenario-based simulation testing is commonly used [25, 27, 28, 34, 37, 38], with a specific focus on collision detection. Such methods often generate a limited number of scenarios (often manually created) under fixed settings [27]. Recent approaches




Figure 1: WPgen in the Context of the MAPLE-K loop

have begun exploring techniques, such as machine learning [25, 28], random testing [38], and search approaches [37] for test automation. In contrast, WPgen focuses on generating minor variations in given waypoints (i.e., tests) that should trigger adaptations in AVs. Moreover, WPgen is a real-world application of search-based software engineering (SBSE) [18] in the maritime domain.

To evaluate WPgen, we used the Marine Systems Simulator (MSS) [1], a widely used maritime simulator for the design, development, and testing of marine applications [9]. We chose three AV models (one overwater tanker, and two underwater) and their navigation software that uses waypoint-based navigation. We compared the three variations of WPgen with Random Search (RS) as baseline. Experimental results showed that the effectiveness of WPgen depends on the AV. Moreover, we study characteristics of various unexpected paths generated by WPgen. Such characterization helps determine which adaptations to trigger for which path types. Finally, we present research and practical implications of WPgen.

2 REAL-WORLD CONTEXT

RoboSAPIENS [2] is a large European initiative focused on developing self-adaptive software for robotic applications across several domains, including maritime, manufacturing, and service. The project aims to develop novel methods for testing and implementing robotic software adaptations in unforeseen situations, enabling robots to handle such situations during real-world operations. The project collaborates with research organizations and industrial partners, providing real-world robotic software use cases for evaluating the developed methods. This paper focuses on our use case from the maritime domain, i.e., AVs. These AVs operate in the marine environment and are subject to rapid changes, e.g., in wind and ocean currents, resulting in uncertain AV behavior that can lead to unsafe situations.

The long-term goal of the project is to implement self-adaptation using the MAPLE-K [23] (see Fig. 1) loop—an extension of the well-established MAPE-K [4] loop—to enable AVs to adapt their behavior in the presence of uncertainties and prevent potentially unsafe situations. A typical MAPE-K loop works as follows: the Monitor component continuously monitors the state of the system (e.g., an AV), which is then analyzed by the Analyze component. If the analysis indicates that the system requires adaptation, the Plan component devises a plan, which is subsequently executed by the Execute component to perform the adaptation on the system. MAPLE-K extends MAPE-K by introducing the Legitimate component, which verifies whether the adaptation meets all safety requirements. All components contribute to the Knowledge component, where relevant information is stored. At the time of development

of the self-adaptive software, it is essential to find settings that should trigger adaptations. This is needed to ensure that not only the triggered adaptations are correct but also that the adaptations’ implementations are correct, preventing the AVs from going into unsafe situations.

The project is currently in its early phase, focusing on implementing the MAPLE-K loop in its robotic software, while our work, in parallel, aims to develop an approach to test the MAPLE-K loop. Specifically, it generates modified waypoints that should trigger relevant adaptations, thereby assessing whether the Plan and Legitimate components can successfully select and verify appropriate adaptations, respectively.

3 APPROACH

3.1 Approach Overview

We represent an AV and its associated *navigation software* as \mathcal{V} . In particular, we consider *waypoint-based* navigation software that guides the AV through a set of predefined *waypoints*. A *waypoint* w_p is a coordinated destination point. For an overwater AV, a waypoint is defined in a North-East coordinate system with two dimensions (i.e., $w_p = (x_i, y_i)$). Similarly, for an underwater AV, the waypoint is defined in a North-East-Down coordinate system with 3 dimensions (i.e., $w_p = (x_i, y_i, z_i)$). A path is defined as a set of N waypoints $\mathcal{WP} = \{w_{p_1}, \dots, w_{p_N}\}$ that the AV \mathcal{V} must navigate through.

Waypoints placement must ensure that the AV has sufficient space to navigate smoothly between them. Specifically, each pair of consecutive waypoints w_{p_i} and $w_{p_{i+1}}$ cannot be closer than $minWPdist$. The value of $minWPdist$ depends on the specific AV.

Our approach WPgen generates waypoints representing possible alterations of the original set of waypoints $\mathcal{WP}_o = \{w_{p_1}^o, \dots, w_{p_N}^o\}$ that the AV \mathcal{V} is originally supposed to follow. These variations could occur during navigation due to environmental disturbances such as wind and sea current, or obstacles that block the planned path.

The goal of WPgen is to find a set of modified waypoints that are challenging for the current navigation software \mathcal{V} , i.e., that lead to obtaining *unstable* paths of the AV. These sets of waypoints will be used to implement a self-adaptive navigation software that can, at runtime, safely self-adapt to these runtime variations. Note that implementing such self-adaptive navigation software is not part of this work. This implementation is being carried out by our industrial partners in RoboSAPIENS, in parallel with this work.

WPgen is cast as a multi-objective optimization problem that is described in detail in the following sections.

3.2 Solution Encoding

In WPgen, the *search variables* \bar{k} represent a set of $N - 1$ consecutive waypoints that define the modified set of waypoints the AV \mathcal{V} must follow during navigation. We do not search for the first waypoint that is the same as that of \mathcal{WP}_o (i.e., $w_{p_1}^o$), as it is the vessel’s current starting point and it must be common to all the \mathcal{V} ’s paths.

For each searched waypoint, there are two or three variables, depending on whether the search is conducted for an overwater or an underwater AV. For simplicity, in the following, we consider an overwater AV (the formalization for an underwater AV is similar);

Algorithm 1: WPgen_{seed} and WPgen_{comb} – Generation of initial individual

Input: $\mathcal{WP}_o = \{wp_1^o, \dots, wp_N^o\}$ (original waypoints)
Output: $\mathcal{WP}' = \{wp_1', \dots, wp_N'\}$ (modified waypoints)

```

1 numOfMutParams  $\leftarrow$  0;
2  $\mathcal{WP}' \leftarrow \mathcal{WP}_o$ ;
3 do
4    $mp \leftarrow \text{randomNumber}(2 \text{ to } N)$ ; //select waypoint  $wp_{mp}'$ 
5    $wp_{mp}' \leftarrow \begin{pmatrix} \text{randomNumber}((v_{x,i}^o - \Delta) \text{ to } (v_{x,i}^o + \Delta)) \\ \text{randomNumber}((v_{y,i}^o - \Delta) \text{ to } (v_{y,i}^o + \Delta)) \end{pmatrix}$ ;
6   numOfMutParams  $\leftarrow$  numOfMutParams + 1;
7    $p \leftarrow \text{rand}()$ ; //returns random value between 0 and 1
8 while  $p < 0.5^{\text{numOfMutParams}}$ ;
9 return  $\mathcal{WP}'$ ;
```

so, the search variables are $\bar{k} = \{k_{2,x}, k_{2,y}, \dots, k_{N,x}, k_{N,y}\}$, where $(k_{i,x}, k_{i,y})$ defines the i th waypoint.

A *solution* (or *individual*) is a concrete assignment $\bar{v} = \{v_{x,2}, v_{y,2}, \dots, v_{x,N}, v_{y,N}\}$, that defines the modified waypoints $wp_2 = (v_{x,2}, v_{y,2}), \dots, wp_N = (v_{x,N}, v_{y,N})$. So, together with the starting waypoint wp_1^o , these define the set of waypoints $\mathcal{WP} = \{wp_1^o, wp_2, \dots, wp_N\}$ that will be executed in the AV's simulator MSS as test for the navigation software.

Since the generated set of waypoints \mathcal{WP} must be similar to the original set of waypoints $\mathcal{WP}_o = \{wp_1^o, \dots, wp_N^o\}$, the search space of variables \bar{k} is defined in the neighborhood of the values of $\mathcal{WP}_o = \{(v_{x,1}^o, v_{y,1}^o), \dots, (v_{x,N}^o, v_{y,N}^o)\}$. Specifically, for each pair of variables $k_{i,x}, k_{i,y}$ (with $i \in \{2, \dots, N\}$), their search spaces are $[v_{x,i}^o - \Delta, v_{x,i}^o + \Delta]$ and $[v_{y,i}^o - \Delta, v_{y,i}^o + \Delta]$, where Δ is a hyperparameter of the approach.

As explained in Sect. 3.1, consecutive waypoints must be at least at a distance of minWPdist . Solutions generated by WPgen could violate such constraint. Therefore, for each individual, WPgen checks if the constraint is satisfied and, if it is not, it discards the individual (i.e., following a death penalty approach [22]).

3.3 Initial population

At the beginning of the search, WPgen must be initialized with an initial population. Prior works have shown that seeding the initial population is beneficial in terms of both effectiveness as well as efficiency of the algorithm when generating and selecting test cases [6, 14, 29]. Because of this, WPgen can be configured to use two different seeding strategies and a traditional fully random strategy, described in the following.

3.3.1 WPgen_{seed} – Closeness seed. This first strategy takes as input the original set of waypoints \mathcal{WP}_o and applies minor mutations by following Algorithm 1. Specifically, the strategy aims at mutating (in the allowed search space) at least one randomly selected waypoint of the initial solution, as shown in Lines 4-5. As the number of mutated variables (i.e., numOfMutParams) increases, the chances for mutating a new variable decreases.

3.3.2 WPgen_{comb} – Combined. The second strategy takes as input the original waypoints \mathcal{WP}_o used for testing the AV and applies minor mutations (following Algorithm 1) to half of the solutions

in the population. The other half of the population is generated entirely randomly (still in the search space of search variables), promoting, in this way, diversity in the population.

3.3.3 WPgen_{rnd} – Random. The last strategy follows the traditionally used fully random generation of the entire population. Since we want to produce waypoints close to the set of original waypoints \mathcal{WP}_o , one of the individuals in the population is \mathcal{WP}_o .

3.4 Objective Functions

Given a set of waypoints \mathcal{WP} obtained from an individual \bar{v} (see Sect. 3.2), the approach simulates the navigation of AV \mathcal{V} over \mathcal{WP} , and obtains the path p followed by \mathcal{V} , i.e.,

$$p = \{(x_1^p, y_1^p), \dots, (x_K^p, y_K^p)\} = \text{simulate}(\mathcal{V}, \mathcal{WP}) \quad (1)$$

If the AV managed to follow all the waypoints in \mathcal{WP} , p can be split in $N - 1$ *sub-paths* sp_1, \dots, sp_{N-1} , i.e., the sections of p between each consecutive pair of waypoints. Note that the AV may fail to follow all the waypoints: in that case, the number of sub-paths is lower than $N - 1$.

WPgen has two competing objectives that capture the overall goal of the approach, i.e., finding the minimal variations of the original waypoints \mathcal{WP}_o that are challenging for the AV.

The first objective aims at minimizing the distance between $\mathcal{WP}_o = \{(v_{x,1}^o, v_{y,1}^o), \dots, (v_{x,N}^o, v_{y,N}^o)\}$ and the waypoints $\mathcal{WP} = \{(v_{x,1}, v_{y,1}), \dots, (v_{x,N}, v_{y,N})\}$ obtained from an individual¹:

$$\text{fit}_{\text{distWPs}}(\mathcal{WP}) = \sqrt{\sum_{i=1}^N ((v_{x,i}^o - v_{x,i})^2 + (v_{y,i}^o - v_{y,i})^2)} \quad (2)$$

The second objective aims at finding *unstable paths*. Intuitively, a path is more unstable if it largely deviates from the optimal path composed of stable segments between consecutive waypoints. So, the fitness function that must be maximized is defined as follows:

$$\text{fit}_{\text{unstable}}(\mathcal{WP}) = \sum_{i=1}^{N-1} \frac{|sp_i|}{\sqrt{(v_{x,i+1}^o - v_{x,i}^o)^2 + (v_{y,i+1}^o - v_{y,i}^o)^2}} \quad (3)$$

where $p = \{sp_1, \dots, sp_{N-1}\} = \text{simulate}(\mathcal{V}, \mathcal{WP})$. The intuition is that the length of each sub-path sp_i (approximated by the number of its points) is normalized by the Euclidean distance between $wp_i^o = (v_{x,i}^o, v_{y,i}^o)$ and $wp_{i+1}^o = (v_{x,i+1}^o, v_{y,i+1}^o)$: the higher this value, the more unstable the sub-path is. The fitness is given by the sum of the instability value of each sub-path.

3.5 Search Operators

WPgen uses classical simulated binary crossover and polynomial mutation operators with the default settings from PlatEMO [33]–the framework we used for the implementation of NSGA-II while adapting them for our specific problem's constraints.

In the crossover operation that recombines two individuals, we guarantee that the randomly selected crossover point occurs only between waypoints. This avoids splitting a waypoint, which could otherwise introduce too big changes to the waypoints.

¹With an abuse of notation, we compute the fitness on the set of waypoints \mathcal{WP} obtained from an individual \bar{v} .

The mutation operator performs changes at a coordinate level. It randomly changes the coordinate with a specified default probability from PlatEMO. While performing mutation, we check whether the constraint given by the solution encoding (see Sect. 3.2) is satisfied, i.e., consecutive waypoints must be at least at a distance of $\min WPdist$; otherwise, the mutation is not applied.

4 EMPIRICAL EVALUATION DESIGN

4.1 Research Questions

We aim to study the effectiveness of the WPgen (with its three different seeding strategies) compared to a simple baseline, i.e., Random Search (RS). Moreover, we aim to study the characteristics of unexpected paths followed by AVs in response to waypoints generated by WPgen. Consequently, we would like to answer the following research questions (RQs):

RQ1: How effective is WPgen compared to random search (RS)? This RQ studies whether the problem we study is complex enough to require the sophisticated multi-objective algorithm NSGA-II or if it can be solved with a simple baseline technique, i.e., RS. To this end, we compared RS with the three versions of WPgen (i.e., for the three seeding strategies) with three different case studies. To perform the analysis, we performed a pairwise comparison between RS and each of the WPgen variations utilizing statistical tests.

RQ2: How do different seeding strategies affect the effectiveness of the WPgen? With this RQ, we aim to answer how the different seeding strategies affect WPgen’s effectiveness. To answer this RQ, we compared the effectiveness of the three seeding strategies, $WPgen_{seed}$, $WPgen_{comb}$ and $WPgen_{rnd}$, for each case study by utilizing statistical tests.

RQ3: How good is WPgen in generating paths leading to different types of unexpected paths? This RQ studies the different path types generated by the three versions of WPgen. Studying this is important for understanding the characteristics of different unexpected paths, as such characterization helps determine which adaptations are required for each path type. To answer this RQ, we classified each case study and search approach’s sub-paths to get insight into the types of paths generated and differences in the different case studies.

4.2 Setup and Execution

4.2.1 Case studies. To evaluate WPgen, we selected the Marine Systems Simulator (MSS) toolbox, which is designed in Matlab/Simulink to implement mathematical models of marine systems [1, 26]. MSS integrates various toolboxes to provide a comprehensive platform simulating different vessels. The simulator enables users to simulate scenarios such as dynamic positioning, autopilot design, and other marine operations. The simulator modular structure allows for flexibility in implementation, making it a practical tool for testing and rapid prototyping AVs, given that doing so on a real AV is expensive. MSS has been used in previous works to simulate ships, generating data for analysis [7], or testing marine operations [35].

Our case studies are three AVs covering underwater and overwater AVs with different environments and a set of initial waypoints WP_o . Specifically, the selected AVs are the waypoint-guided AVs in MSS and comprise of one overwater tanker system, *Mariner*, and two underwater AVs, *Remus100* and *Nspauv*. The number of

waypoints in WP_o (i.e., N) is 6 for *Mariner*, and 7 for *Remus100* and *Nspauv*, provided by the MSS initial configuration of the WP .

4.2.2 Search algorithms. We implemented WPgen using the evolutionary multi-objective optimization platform PlatEMO [33]. We utilized their implementation of NSGA-II [11] and adapted their crossover operator to our operator described in Sect. 3.5. Moreover, we replaced the default population generation of PlatEMO with the three different seeding strategies inside NSGA-II, resulting in three variations of WPgen. As explained in Sect. 3.3, the specific seeding approaches are (i) a fully seeded initial population ($WPgen_{seed}$), (ii) a combination of seeded and randomly generated waypoints ($WPgen_{comb}$), and (iii) a fully randomly generated initial population ($WPgen_{rnd}$). To guide the search to WP close to WP_o , all the initial populations of the different approaches contain one individual that is equal to WP_o .

For our baseline algorithm, Random Search (RS), we generate the same number of individuals as for WPgen such that the total number of fitness evaluations is equal for WPgen and RS to ensure a fair comparison. The initial population for RS is also generated randomly, as discussed in Sect. 3.2.

4.2.3 Parameter Settings. By following the guide of Arcuri and Briand [5], each variation of WPgen and RS was run 30 times to account for the randomness of the search algorithms. Given that there are four approaches (the three seeding strategies of WPgen and RS), in total we conducted 120 runs for each case study.

Due to the long execution time of the MSS simulator, an initial preliminary experiment was conducted to choose the most appropriate parameter settings, which involved analyzing the convergence of the Pareto Front. The population size was set to 10, and the maximum number of generations to 1000, resulting in a total number of 10,000 fitness evaluations. The default mutation operator and our slightly adapted crossover operator (see Sect. 3.5) were used with default settings from PlatEMO. Specifically, the probability of crossover and the expectation of the number of mutated variables were both set to 1, whereas the distribution index of crossover and mutation were both set to 20.

The hyperparameter Δ that defines the search space (see Sect. 3.2) was set to 400 and 150 for the overwater AV and underwater AVs, respectively. These values were set based on the AVs’ characteristics, such as AV length and original distance between waypoints.

4.3 Evaluation Metrics

4.3.1 Quality indicator. To assess the effectiveness of each approach, we employed the *hypervolume* (HV) metric [31] as our quality indicator. HV measures the volume of the objective space covered by the Pareto front with respect to a reference point, which is chosen based on the worst value observed across all approaches and experiments. A higher HV value indicates better coverage and, thus, better performance and diversity of an approach.

4.3.2 Statistical tests. To compare the performance of the search approaches in RQ1 and RQ2, we employed the Mann-Whitney U-test for statistical significance and Vargha Delaney \hat{A}_{12} to measure the effect size following the guide [5]. Specifically, we compared each variation of WPgen among themselves and with RS, resulting in a total of six comparisons for each AV.

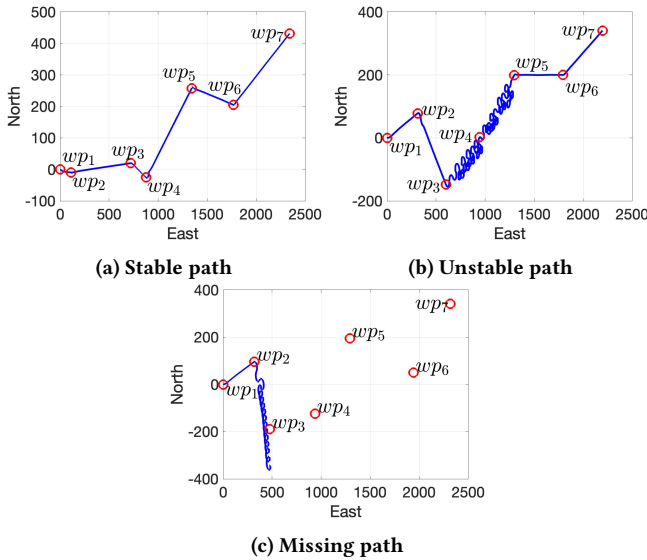


Figure 2: Example of paths of Remus100 (North-East view) with different sub-path types: stable, unstable, and missing. A waypoint is shown as a red circle and labeled as wp_i for i th waypoint. A path consists of waypoints from wp_1 to wp_7 . Fig. 2b shows the AV’s path is unstable between some waypoints. Fig. 2c shows the AV missing several waypoints.

The Mann-Whitney U-test gives a p-value that assesses whether there is a statistical significant difference between two approaches A and B in terms of HV. If the p-value is above the significance level α ($\alpha = 0.05$ in this work), it means there is *No statistically significant Difference* (ND). Instead, if the p-value is below α , there is a statistically significant difference. In this case, we check the effect size \hat{A}_{12} that tells the probability that one algorithm outperforms the other: if $\hat{A}_{12} > 0.5$, it means that A is significantly better than B ; otherwise, if $\hat{A}_{12} < 0.5$, B is significantly better than A . We also check the strength of the effect size following the classification in [24]: negligible when $\hat{A}_{12} \in (0.44, 0.56)$, small when $\hat{A}_{12} \in (0.34, 0.44)$ or $\hat{A}_{12} \in [0.56, 0.64)$, medium when $\hat{A}_{12} \in (0.29, 0.34)$ or $\hat{A}_{12} \in [0.64, 0.71)$, and large when $\hat{A}_{12} \in [0, 0.29)$ or $\hat{A}_{12} \in [0.71, 1]$.

4.3.3 Classification of sub-paths and paths. For RQ3, we are interested in studying the types of generated paths. We have defined three categories, *stable*, *unstable*, and *missing*. The three sub-path types are shown in Fig. 2. The stability/unstability of an AV can be observed by measuring how the angles of the AV change, such as rolling from side to side. To detect this unstable behavior, we calculate the autocorrelation of the axis measurements to quantify the instability and differentiate between a stable measurement and an unstable one by calculating the number of peaks and the magnitude of each peak. We counted one angle as unstable if it had more than one threshold peak above 0.1 [30]. With this, we can differentiate between a stable and an unstable sub-path based on how many axes are stable. Specifically: (1) *stable*: all of the angles are stable, (2) *unstable*: at least one of the angles is unstable, and (3) *missing*: the vessel fails to reach a waypoint.

Table 1: RQ1 and RQ2 – Results of the statistical comparison between $WPgen_{seed}$, $WPgen_{comb}$, $WPgen_{rnd}$, and RS, in terms of HV. ND indicates no significant difference (i.e., p-value ≥ 0.05). The *Best Approach* column identifies the approach that is significantly better than the other (i.e., p-value < 0.05). The *Strength* column indicates the strength of the significance based on the \hat{A}_{12} effect size, as described in Sect. 4.3.2

AV \mathcal{V}	Compared approaches		Best approach	Strength
	A	B		
Mariner	RS	$WPgen_{seed}$	$WPgen_{seed}$	large
	RS	$WPgen_{comb}$	$WPgen_{comb}$	large
	RS	$WPgen_{rnd}$	$WPgen_{rnd}$	large
	$WPgen_{seed}$	$WPgen_{comb}$	ND	ND
	$WPgen_{seed}$	$WPgen_{rnd}$	ND	ND
	$WPgen_{comb}$	$WPgen_{rnd}$	ND	ND
Remus100	RS	$WPgen_{seed}$	RS	large
	RS	$WPgen_{comb}$	$WPgen_{comb}$	large
	RS	$WPgen_{rnd}$	$WPgen_{rnd}$	large
	$WPgen_{seed}$	$WPgen_{comb}$	$WPgen_{comb}$	large
	$WPgen_{seed}$	$WPgen_{rnd}$	$WPgen_{rnd}$	large
	$WPgen_{comb}$	$WPgen_{rnd}$	ND	ND
Nspauv	RS	$WPgen_{seed}$	ND	ND
	RS	$WPgen_{comb}$	ND	ND
	RS	$WPgen_{rnd}$	ND	ND
	$WPgen_{seed}$	$WPgen_{comb}$	ND	ND
	$WPgen_{seed}$	$WPgen_{rnd}$	ND	ND
	$WPgen_{comb}$	$WPgen_{rnd}$	ND	ND

Using the previous classification sub-paths, we have a metric for describing the full path. Specifically, given the $N - 1$ sub-paths of a path, a full path can have $2^{N-1} + 2(N - 2)$ different categorizations.

5 RESULTS AND ANALYSES

5.1 RQ1 – Comparison with Random Search

Table 1 reports the results of the statistical tests when comparing the three variations of $WPgen$ with RS. For each AV, the comparison with RS is shown in rows where the *Approach A* column has a value of RS. The fourth column indicates which of the two approaches performs significantly better, followed by the strength category computed based on \hat{A}_{12} (see Sect. 4.3.2). If no conclusion can be drawn from the statistical test (i.e., p-value < 0.05), ND is written to indicate no significant difference between the two compared approaches. Moreover, the descriptive statistics of the results are shown as the boxplots showing the distributions of the HV values for each AV and approach in Fig. 3.

For Mariner, when analyzing the distributions of the HV values in Fig. 3a, we can see that RS is worse than the three versions of $WPgen$. Moreover, RS has more distribution variance than $WPgen$, indicating that the effectiveness of RS is unreliable. Furthermore, the results of statistical tests in Table 1 show that each variation of $WPgen$ is significantly better than RS with Large strength.

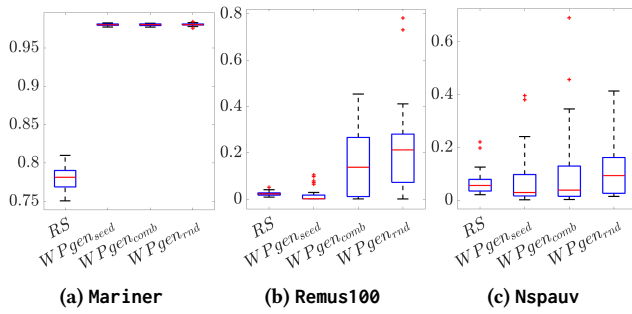


Figure 3: RQ1 and RQ2 – Distribution of the calculated HV metric for each autonomous vessel and approach

For Remus100, Fig. 3b shows that WPgen_{comb} and WPgen_{rnd} perform better than RS; however, the variance in their distribution is high. WPgen_{seed}, on the other hand, seems to perform worse than RS. Both RS and WPgen_{seed} have less variances in their distributions. These observations are confirmed by the results of the statistical tests in Table 1, where we can see that WPgen_{comb} and WPgen_{rnd} perform significantly better than RS with large strength. However, RS performs significantly better than WPgen_{seed} with large strength.

For Nspauv, when comparing the distributions of the HV values in Fig. 3c between RS and each of three variations of WPgen, we do not see clearly which approach performs the best. This is further confirmed by the results of statistical tests reported in Table 1 showing no significant differences among any pair of approaches.

Conclusion for RQ1: WPgen either performs significantly better than RS or shows no significant differences depending on the autonomous vessel, except Remus100, where RS is significantly better than WPgen_{seed}. Thus, we can conclude that our problem is complex enough to require a guided algorithm such as NSGA-II, except Nspauv.

5.2 RQ2 – Comparison of Seeding Strategies

Fig. 3 reports the distribution of the HV values for the three version of WPgen (i.e., WPgen_{seed}, WPgen_{comb}, and WPgen_{rnd}), while Table 1 reports the results of their statistical comparison.

For Mariner, when comparing the distributions of the three variations of WPgen in Fig. 3a, we do not see any differences; this is further confirmed by the results of statistical tests in Table 1, which show no significant differences among the three approaches.

For Remus100, when looking at the distributions of the three variations of WPgen in Fig. 3b, we see no clear difference between WPgen_{comb} and WPgen_{rnd}; this is further confirmed by the results of statistical tests in Table 1, which show no differences between the two approaches. Instead, in Fig. 3b, WPgen_{seed} seems worse than the other two approaches; indeed, results of the statistical tests in Table 1 show that WPgen_{seed} is significantly worse than WPgen_{comb} and WPgen_{rnd} with large strength. Recall from RQ1 (see Sect. 5.1) that WPgen_{seed} was also outperformed by RS for Remus100.

For Nspauv, when looking at the distributions of three approaches in Fig. 3c, we do not see clear differences; this is confirmed by the results in Table 1 that report no significant differences.

Table 2: RQ3 – Percentage (across 30 runs) of different sub-paths (i.e., stable, unstable, and missing) from all the solutions on a Pareto Front

Approach	Type of path	AV \mathcal{V}		
		Mariner	Remus100	Nspauv
WPgen _{seed}	Stable	16.89%	74.72%	79.84%
	Unstable	83.11%	20.17%	18.20%
	Missing	0	5.12%	1.95%
WPgen _{comb}	Stable	17.57%	68.98%	68.88%
	Unstable	82.43%	24.47%	25.40%
	Missing	0	6.54%	5.72%
WPgen _{rnd}	Stable	18.72%	68.12%	69.80%
	Unstable	81.28%	25.23%	25.30%
	Missing	0	6.65%	4.90%

Conclusion for RQ2: The suitability of a seeding strategy depends on the AV \mathcal{V} being tested. For instance, for Mariner and Nspauv, it does not matter which strategy to use, whereas, for Remus100, either WPgen_{comb} or WPgen_{rnd} is the best strategy.

5.3 RQ3 – Paths Classification

Table 2 reports, for each approach and AV, the percentage of different sub-paths calculated from all solutions on the Pareto front (average across 30 runs).

For Mariner, the percentages of stable sub-paths across the three approaches are similar: 16.89%, 17.57%, and 18.72% for WPgen_{seed}, WPgen_{comb}, and WPgen_{rnd}, respectively. Similarly, the percentages for unstable paths are also very close: 83.11%, 82.43%, and 81.28%. None of the three approaches found any missing sub-paths. These results suggest that all three approaches perform similarly in identifying unstable sub-paths for Mariner. Furthermore, the higher percentage of unstable sub-paths indicates that it was relatively easier for any approach to detect unstable paths for Mariner.

For Remus100, WPgen_{seed} identified a slightly higher percentage of stable sub-paths (74.72%) compared to WPgen_{comb} (68.98%) and WPgen_{rnd} (68.12%). In contrast, for unstable sub-paths, WPgen_{seed} had a slightly lower percentage (20.17%) compared to WPgen_{comb} (24.47%) and WPgen_{rnd} (25.23%). Regarding missing sub-paths, all three approaches exhibited similar percentages: 5.1% for WPgen_{seed}, 6.54% for WPgen_{comb}, and 6.65% for WPgen_{rnd}. These results suggest that WPgen_{comb} and WPgen_{rnd} are slightly better than WPgen_{seed} at identifying unstable and missing sub-paths.

For Nspauv, the results are similar to those of Remus100. WPgen_{seed} achieved the highest percentage of stable sub-paths (79.84%) compared to WPgen_{comb} (68.88%) and WPgen_{rnd} (69.80%). WPgen_{comb} (25.4%) and WPgen_{rnd} (25.3%) had higher percentages of unstable sub-paths than WPgen_{seed} (18.2%). Similarly, WPgen_{comb} (5.72%) and WPgen_{rnd} (4.92%) recorded higher percentages of missing sub-paths than WPgen_{seed} (1.95%). Once again, these results suggest the better suitability of WPgen_{comb} and WPgen_{rnd} for identifying unstable and missing sub-paths for Nspauv.

Table 3: RQ3 – Average (across 30 runs) percentage of unique full paths from all the solutions on a Pareto Front. The percentage is based on the maximum number of possible paths

Approach	AV \mathcal{V}		
	Mariner	Remus100	Nspauv
WPgen _{seed}	9.65%	23.24%	36.35%
WPgen _{comb}	12.00%	30.45%	70.62%
WPgen _{rnd}	13.15%	37.88%	74.81%

As explained in Sect. 4.3.3, each path followed by a vessel can be classified by considering the combination of the classifications of its sub-paths; at most, there are $2^{N-1} + 2(N-2)$ possible types of paths. Table 3 indicates the average percentage of unique full paths generated per approach for each AV across the 30 runs. For Mariner, WPgen_{rnd} manages to generate, on average, slightly more unique paths (i.e., 13.15%) than the other approaches, i.e., 12.00% for WPgen_{comb} and 9.65% for WPgen_{seed}. A similar trend is observed for Remus100 and Nspauv, where the average percentage of unique paths is highest for WPgen_{rnd} and lowest for WPgen_{seed}. Such results are expected since total random seeding will result in more unique paths, whereas fully seeding the initial population causes the approach to focus on those waypoints. At the same time, the combination of both approaches produces results that fall in the middle. Another observation is that Mariner, generally across all approaches, has the lowest average percentage of unique paths, followed by Remus100, whereas Nspauv has the highest average number of unique paths. This result suggests that the number of unique paths will vary depending on the AV.

Conclusion for RQ3: For underwater autonomous vessels (Remus100 and Nspauv), WPgen_{comb} and WPgen_{rnd}, which incorporate random individuals in the initial population, generate more paths with unstable and missing sub-paths. For overwater autonomous vessels (Mariner), all approaches perform similarly in identifying unstable and missing sub-paths. Overall, for any autonomous vessel type, approaches with random initial populations are recommended. Moreover, WPgen_{seed} identifies fewer unique paths than WPgen_{comb} and WPgen_{rnd}, with the number of unique paths varying across vessels.

6 DISCUSSION

6.1 Research and Practical Implications

Based on the results from RQ1, we found that all three variations of WPgen either perform significantly better than RS or show no significant differences, depending on the AV. Similarly, results for RQ2 showed differences among the approaches for different AVs. These results suggest that the selection of a particular WPgen depends on the AV being tested. For Mariner, while there were no significant differences among the three variations of WPgen, all were significantly better than RS, indicating that any variation of WPgen is sufficient for testing Mariner. For Remus100, WPgen proved to be the best, making it the natural choice for testing. For Nspauv, there were no significant differences between RS and the three variations

of WPgen, meaning that even RS is sufficient for testing Nspauv. Although we now have preliminary indications for selecting approaches, it raises further questions about why a specific approach is better for a particular AV. This requires additional empirical evaluations that consider more AV characteristics. Such a study is part of our future work.

From RQ3, we found that approaches incorporating randomness in the initial population outperformed those that fully seed the initial population for both underwater and overwater AVs. Therefore, such approaches—WPgen_{comb} and WPgen_{rnd}—are recommended for identifying waypoints that should trigger adaptations in self-adaptive AVs. From RQ3, we also found that WPgen_{seed} generally identified fewer unique full paths than WPgen_{comb} and WPgen_{rnd}. These results were expected, as WPgen_{seed} is designed to enable the search algorithm to focus around the seeded waypoints. Depending on the context, a tester may want to focus only on a given set of waypoints; in that case, WPgen_{seed} is recommended. However, it naturally lacks the diverse coverage the other two approaches provide. On the other hand, if a tester seeks more diverse paths, WPgen_{comb} or WPgen_{rnd} are recommended.

In general, we conclude that WPgen can find unstable paths close to the original trajectory \mathcal{WP}_o . Thus, our results bring new knowledge about the importance of \mathcal{WP} generation and analysis of AVs. Furthermore, the waypoints that result in unstable paths can be utilized for different purposes. First, these waypoints can be used to trigger different adaptations in self-adaptive AVs to verify whether correct adaptations are triggered and implemented. Second, further analysis can be conducted to identify any patterns that explain why these \mathcal{WP} specifically result in unstable paths and how such \mathcal{WP} can be avoided in the future. Finally, our methods for categorizing the sub-paths can be further utilized to create more advanced fitness functions to further enhance the performance of WPgen.

Our approach WPgen can be used to pre-evaluate a set of waypoints \mathcal{WP} (constructed as minor variations of the original ones \mathcal{WP}_o), and determine the safety of the AV. Detection of a probable unstable path can be utilized as feedback to the system, letting the AV adjust its controllers to get a smoother and safer path. A set of waypoints \mathcal{WP} that lead to a path having unstable or missing sub-paths, can be utilized for different AV navigation software to compare and determine if it is better at handling \mathcal{WP} . Navigation software can also utilize the characteristics of the \mathcal{WP} to recognize critical situations that require adaptation.

6.2 Threats to Validity

Internal validity. An internal validity threat is related to the parameters of the selected algorithm, i.e., NSGA-II, which is the core part of the three variations of WPgen. To mitigate this threat, we used the default crossover and mutation operators with default parameter settings from PlatEMO. These parameters were then kept the same for the three variations of WPgen for a fair comparison. One could argue that we selected only one algorithm, i.e., NSGA-II; however, it is widely used in search-based software engineering, and we naturally plan to investigate additional algorithms in the future.

External validity We experimented with three AVs of different characteristics in a fixed environment. Thus, the results may not generalize to other AVs or environments. This highlights the need

for additional empirical evaluations, which, along with exploring other search algorithms, will form the basis of our future studies.

Conclusion validity All search algorithms have inherent randomness. To address this, we executed each approach and RS 30 times, following a well-established guide [5]. Additionally, to ensure greater confidence in the results, we analyzed them using well-established statistical tests as recommended in the same guide.

Construct validity A concern is related to measures used to draw conclusions. For example, to automatically classify paths, we used autocorrelation; however, other metrics, such as zero-crossing or frequency analysis, also exist. We did not select these metrics because they were less precise and misclassified some stable paths as unstable during our exploratory study. Another concern relates to the stopping criteria, which in our context was defined as the number of fitness evaluations. Based on an exploratory study, this was set to 10,000 and kept consistent across all approaches.

7 RELATED WORK

In this work, we proposed an approach to generate tests exposing cases in which an AV follows unstable paths. Such cases will be considered during the development of a self-adaptive AV in the context of our project (see Sect. 2). To the best of our knowledge, this is the first approach proposed for this goal. On the other hand, in the self-adaptive community [8, 10], some works have been proposed to test the self-adaptive system itself. H nsel et al. [17] proposed a model-based testing approach for MAPE-K that relies on architectural runtime models (RTM). Specifically, the approach checks at runtime the conformance between the running self-adaptive system and the oracle RTM. Eberhardinger et al. [12] proposed an approach based on model checking to generate test cases for self-adaptive systems. In order to make the approach scalable, they simplify the system using the Corridor Enforcing Infrastructure (CEI) architectural pattern. Arcaini et al. [3] also proposed an approach based on model checking to derive test cases from self-adaptive ASMs, a formal notation for the analysis of self-adaptive systems. Fredericks et al. [15, 16], instead, proposed online testing techniques able to better adjust to the uncertainty of the environment.

Our work is also related to Search-Based Software Engineering (SBSE), which formulates software engineering problems as optimization problems and solves them using search algorithms, such as single-objective and multi-objective search algorithms [18]. Since its inception, SBSE has been applied to various software engineering sub-areas, e.g., requirements engineering, modeling, testing, and debugging [19]. Within SBSE, Search-Based Testing (SBT) focuses specifically on testing software systems using search algorithms, which is the focus of this paper. SBT has been applied to testing a variety of systems, including autonomous driving systems [32], autonomous robots [20], unmanned aerial vehicles (UAVs) [21], automated warehouses [36], and sports-related cyber-physical systems [36]. Despite the applications of SBT in various domains, its application for testing the navigation software of AVs remains understudied. To this end, this paper presents an application of SBT to generate tests (i.e., minor variations in the predefined waypoints) that should trigger adaptations in AVs.

Scenario-based testing with simulators is one of the key methods for testing AVs at design time [25, 27, 28, 34, 37, 38], which is also

related to WPgen. Most scenario-based testing works have primarily focused on collision avoidance. Furthermore, existing works often generate a limited number of scenarios with fixed settings, and such generation is often manual, as concluded in a literature review by Porres et al. [27]. Nonetheless, the need for scenario-based testing is being recognized, and recent works have explored more systematic and automated approaches to simulation-based testing for collision avoidance, such as using machine learning techniques [25, 28], random testing [38], or a search-based approach [37]. Our work distinguishes itself from existing works by generating tests that should trigger adaptation in AVs. Specifically, we generate variations of predefined waypoints with SBT, so reducing the manual effort required to identify variations that could potentially impact the AV's ability to follow its intended path.

8 CONCLUSIONS AND FUTURE WORK

We presented a multi-objective search approach, WPgen, to generate tests that should trigger adaptations in maritime autonomous vessels (AVs), where each test is a slight variation of the pre-defined waypoints that the AV must follow during navigation. The approach was further varied into three versions, each employing a different seeding strategy to study how initial population seeding affects performance. We compared the three variations with RS and with each other. Based on the results, we found that, in general, WPgen either performs significantly better than RS or similarly. Moreover, the best-performing variation of the approach depends on the AV. In the future, we will extend WPgen from several perspectives, including defining domain-specific mutation and crossover operators. Moreover, the empirical evaluation will be extended with more algorithms, different environmental setups, and similar systems, such as unmanned aerial vehicles. Finally, we also want to explore machine learning algorithms to detect unstable paths, gain insights into the different characteristics of an unstable path, and determine the root cause of the unstable paths.

Replication package: For replicability, we make all our code, experimental scripts, results, and detailed plots available at Zenodo <https://zenodo.org/records/15093003>.

ACKNOWLEDGMENTS

K. Nyl nder and S. Ali are supported by the Co-tester project (Project# 314544), funded by the Research Council of Norway. S. Ali is also supported by the RoboSAPIENS project, funded by the European Commission's Horizon Europe programme (Grant Agreement # 101133807). P. Arcaini is supported by Engineerable AI Techniques for Practical Applications of High-Quality Machine Learning-based Systems Project (Grant Number JPMJMI20B8), JST-Mirai. Aitor Arrieta is part of the Systems and Software Engineering group of Mondragon Unibertsitatea (IT1519-22), supported by the Department of Education, Universities and Research of the Basque Country.

REFERENCES

- [1] [n. d.]. Marine Systems Simulator. <https://www.fossen.biz/MSS/>. [Online; accessed 25-March-2025].
- [2] [n. d.]. RoboSAPIENS. <https://robosapiens-eu.tech/>. [Online; accessed 25-March-2025].
- [3] Paolo Arcaini, Raffaella Mirandola, Elvinia Riccobene, and Patrizia Scandurra. 2020. Model-Based Testing for MAPE-K adaptation control loops. In *2020 IEEE*

- International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. 43–51. <https://doi.org/10.1109/ICSTW50294.2020.00024>
- [4] Paolo Arcaini, Elvinia Riccobene, and Patrizia Scandurra. 2015. Modeling and Analyzing MAPE-K Feedback Loops for Self-Adaptation. In *Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems* (Florence, Italy) (SEAMS '15). IEEE Press, 13–23. <https://doi.org/10.1109/SEAMS.2015.10>
 - [5] Andrea Arcuri and Lionel Briand. 2011. A Practical Guide for Using Statistical Tests to Assess Randomized Algorithms in Software Engineering. In *Proceedings of the 33rd International Conference on Software Engineering* (Waikiki, Honolulu, HI, USA) (ICSE '11). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1985793.1985795>
 - [6] Aitor Arrieta, Pablo Valle, Joseba A. Agirre, and Goiuria Sagardui. 2023. Some Seeds Are Strong: Seeding Strategies for Search-based Test Case Selection. *ACM Trans. Softw. Eng. Methodol.* 32, 1, Article 17 (Feb. 2023), 47 pages. <https://doi.org/10.1145/3532182>
 - [7] Alexandra Baier and Steffen Staab. 2022. A Simulated 4-DOF Ship Motion Dataset for System Identification under Environmental Disturbances. <https://doi.org/10.18419/DARUS-2905>
 - [8] Betty H. Cheng, Rogério Lemos, Holger Giese, Paola Inverardi, Jeff Magee, Jesper Andersson, Basil Becker, Nelly Bencomo, Yuriy Brun, Bojan Cucik, Giovanna Marzo Serugendo, Schahram Dustdar, Anthony Finkelstein, Cristina Gacek, Kurt Geihs, Vincenzo Grassi, Gabor Karsai, Holger M. Kienle, Jeff Kramer, Marin Litoiu, Sam Malek, Raffaella Mirandola, Hausi A. Müller, Sooyong Park, Mary Shaw, Matthias Tichy, Massimo Tivoli, Danny Weyns, and Jon Whittle. 2009. Software Engineering for Self-Adaptive Systems. Springer-Verlag, Berlin, Heidelberg, Chapter Software Engineering for Self-Adaptive Systems: A Research Roadmap, 1–26. https://doi.org/10.1007/978-3-642-02161-9_1
 - [9] Krzysztof Czaplewski and Piotr Zwolan. 2016. A Vessel's Mathematical Model and its Real Counterpart: A Comparative Methodology Based on a Real-world Study. *The Journal of Navigation* 69, 6 (2016), 1379–1392.
 - [10] Rogério de Lemos, Holger Giese, Hausi A. Müller, Mary Shaw, Jesper Andersson, Marin Litoiu, Bradley Schmerl, Gabriel Tamura, Norha M. Villegas, Thomas Vogel, Danny Weyns, Luciano Baresi, Basil Becker, Nelly Bencomo, Yuriy Brun, Bojan Cucik, Ron Desmarais, Schahram Dustdar, Gregor Engels, Kurt Geihs, Karl M. Göschka, Alessandra Gorla, Vincenzo Grassi, Paola Inverardi, Gabor Karsai, Jeff Kramer, Antónia Lopes, Jeff Magee, Sam Malek, Serge Mankovskii, Raffaella Mirandola, John Mylopoulos, Oscar Nierstrasz, Mauro Pezzè, Christian Prehofer, Wilhelm Schäfer, Rick Schlichting, Dennis B. Smith, João Pedro Sousa, Ladan Tahvildari, Kenny Wong, and Jochen Wuttke. 2013. Software engineering for self-adaptive systems: A second research roadmap. In *Software Engineering for Self-Adaptive Systems II*, Rogério de Lemos, Holger Giese, Hausi A. Müller, and Mary Shaw (Eds.). Springer, 1–32. https://doi.org/10.1007/978-3-642-35813-5_1
 - [11] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Trans. Evol. Comp.* 6, 2 (April 2002), 182–197. <https://doi.org/10.1109/4235.996017>
 - [12] Benedikt Eberhardinger, Hella Seebach, Alexander Knapp, and Wolfgang Reif. 2014. Towards Testing Self-organizing, Adaptive Systems. In *Proceedings of the 26th IFIP WG 6.1 International Conference on Testing Software and Systems - Volume 8763* (Madrid, Spain) (ICTSS 2014). Springer-Verlag New York, Inc., New York, NY, USA, 180–185. https://doi.org/10.1007/978-3-662-44857-1_13
 - [13] Thor Fossen. 2021. *Handbook of Marine Craft Hydrodynamics and Motion Control*. <https://doi.org/10.1002/9781119994138>
 - [14] Gordon Fraser and Andrea Arcuri. 2012. The Seed is Strong: Seeding Strategies in Search-Based Software Testing. In *Proceedings of the 2012 IEEE Fifth International Conference on Software Testing, Verification and Validation (ICST '12)*. IEEE Computer Society, USA, 121–130. <https://doi.org/10.1109/ICST.2012.92>
 - [15] Erik M. Fredericks, Byron DeVries, and Betty H. C. Cheng. 2014. Towards Run-Time Adaptation of Test Cases for Self-Adaptive Systems in the Face of Uncertainty. In *Proceedings of the 9th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2014)*. Association for Computing Machinery, New York, NY, USA, 17–26. <https://doi.org/10.1145/2593929.2593937>
 - [16] Erik M. Fredericks, Andres J. Ramirez, and Betty H. C. Cheng. 2013. Towards runtime testing of dynamic adaptive systems. In *Proceedings of the 8th International Symposium on Software Engineering for Adaptive and Self-Managing Systems* (San Francisco, California) (SEAMS '13). IEEE Press, 169–174.
 - [17] Joachim Hänsel, Thomas Vogel, and Holger Giese. 2015. A Testing Scheme for Self-Adaptive Software Systems with Architectural Runtime Models. In *2015 IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops*. 134–139. <https://doi.org/10.1109/SASOW.2015.27>
 - [18] Mark Harman and Bryan F Jones. 2001. Search-based software engineering. *Information and Software Technology* 43, 14 (2001), 833–839. [https://doi.org/10.1016/S0950-5849\(01\)00189-6](https://doi.org/10.1016/S0950-5849(01)00189-6)
 - [19] Mark Harman, S. Afshin Mansouri, and Yuanyuan Zhang. 2012. Search-based software engineering: Trends, techniques and applications. *ACM Comput. Surv.* 45, 1, Article 11 (Dec. 2012), 61 pages. <https://doi.org/10.1145/2379776.2379787>
 - [20] Dmytro Humeniuk, Foutse Khomh, and Giuliano Antoniol. 2023. AmbieGen: A search-based framework for autonomous systems testing. *Science of Computer Programming* 230 (2023), 102990. <https://doi.org/10.1016/j.scico.2023.102990>
 - [21] Sajad Khatiri, Sebastiano Panichella, and Paolo Tonella. 2023. Simulation-based Test Case Generation for Unmanned Aerial Vehicles in the Neighborhood of Real Flights. In *2023 IEEE Conference on Software Testing, Verification and Validation (ICST)*. 281–292. <https://doi.org/10.1109/ICST57152.2023.00034>
 - [22] Oliver Kramer. 2010. A review of constraint-handling techniques for evolution strategies. *Appl. Comp. Intell. Soft Comput.* 2010, Article 3 (Jan. 2010), 19 pages. <https://doi.org/10.1155/2010/185063>
 - [23] Peter G Larsen, Shaukat Ali, Roland Behrens, Ana Cavalcanti, Claudio Gomes, Guoyuan Li, Paul De Meulenaere, Mikkel L Olsen, Nikolaos Passalis, Thomas Peyrucain, et al. 2024. Robotic safe adaptation in unprecedented situations: the RoboSAPIENS project. *Research Directions: Cyber-Physical Systems 2* (2024), e4.
 - [24] Salvatore S Mangiafico. 2016, revised 2024. Summary and Analysis of Extension Program Evaluation in R. <http://rcompanion.org/handbook/> (Pdf version: <http://rcompanion.org/documents/RHandbookProgramEvaluation.pdf>). Version 1.20.07.
 - [25] Lokukaluge P. Perera. 2019. Deep Learning Toward Autonomous Ship Navigation and Possible COLREGs Failures. *Journal of Offshore Mechanics and Arctic Engineering* 142, 3 (12 2019), 031102. <https://doi.org/10.1115/1.4045372>
 - [26] Tristan Perez, Oyvind Smogeli, Thor Fossen, and Asgeir J Sorensen. 2006. An overview of the marine systems simulator (MSS): A Simulink toolbox for marine control systems. *Modeling, identification and Control* 27, 4 (2006), 259–275.
 - [27] Ivan Porres, Sepinoud Azimi, Sébastien Lafond, Johan Lilius, Johanna Salokannel, and Mirva Salokorpi. 2020. On the Verification and Validation of AI Navigation Algorithms. In *Global Oceans 2020: Singapore – U.S. Gulf Coast*. 1–8. <https://doi.org/10.1109/IEECONF38699.2020.9389133>
 - [28] Ivan Porres, Sepinoud Azimi, and Johan Lilius. 2020. Scenario-based Testing of a Ship Collision Avoidance System. In *2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 545–552. <https://doi.org/10.1109/SEAA51224.2020.00090>
 - [29] José Miguel Rojas, Gordon Fraser, and Andrea Arcuri. 2016. Seeding strategies in search-based unit test generation. *Softw. Test. Verif. Reliab.* 26, 5 (Aug. 2016), 366–401. <https://doi.org/10.1002/stvr.1601>
 - [30] Patrick Schober, Christa Boer, and Lothar A Schwarte. 2018. Correlation coefficients: appropriate use and interpretation. *Anesthesia & analgesia* 126, 5 (2018), 1763–1768.
 - [31] Ke Shang, Hisao Ishibuchi, Linjun He, and Lie Meng Pang. 2021. A Survey on the Hypervolume Indicator in Evolutionary Multiobjective Optimization. *IEEE Transactions on Evolutionary Computation* 25, 1 (2021), 1–20. <https://doi.org/10.1109/TEVC.2020.3013290>
 - [32] Shuncheng Tang, Zhenya Zhang, Yi Zhang, Jixiang Zhou, Yan Guo, Shuang Liu, Shengjian Guo, Yan-Fu Li, Lei Ma, Yinxing Xue, and Yang Liu. 2023. A Survey on Automated Driving System Testing: Landscapes and Trends. *ACM Trans. Softw. Eng. Methodol.* 32, 5, Article 124 (July 2023), 62 pages. <https://doi.org/10.1145/3579642>
 - [33] Ye Tian, Ran Cheng, Xingyi Zhang, and Yaochu Jin. 2017. PlatEMO: A MATLAB platform for evolutionary multi-objective optimization. *IEEE Computational Intelligence Magazine* 12, 4 (2017), 73–87.
 - [34] Tobias Rye Torben, Jon Arne Glomsrud, Tom Arne Pedersen, Ingrid B Utne, and Asgeir J Sørensen. 2023. Automatic simulation-based testing of autonomous ships using Gaussian processes and temporal logic. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability* 237, 2 (2023), 293–313. <https://doi.org/10.1177/1748006X211069277>
 - [35] Tongtong Wang, Guoyuan Li, Baiheng Wu, Vilmar Æsøy, and Houxiang Zhang. 2021. Parameter identification of ship manoeuvring model under disturbance using support vector machine method. *Ships and Offshore Structures* 16, sup1 (2021), 13–21.
 - [36] Man Zhang, Shaukat Ali, and Tao Yue. 2019. Uncertainty-wise test case generation and minimization for Cyber-Physical Systems. *Journal of Systems and Software* 153 (2019), 1–21. <https://doi.org/10.1016/j.jss.2019.03.011>
 - [37] Rong Zhen, Ziqiang Shi, Qiyong Gu, and Shenhua Yang. 2024. A novel deterministic search-based algorithm for multi-ship collaborative collision avoidance decision-making. *Ocean Engineering* 292 (2024), 116524. <https://doi.org/10.1016/j.oceaneng.2023.116524>
 - [38] Feixiang Zhu, Zhengyu Zhou, and Hongrui Lu. 2022. Randomly Testing an Autonomous Collision Avoidance System with Real-World Ship Encounter Scenario from AIS Data. *Journal of Marine Science and Engineering* 10, 11 (2022). <https://doi.org/10.3390/jmse10111588>