Artificial Life Manuscript Submission

Lexicase Selection for Multi-task Evolutionary Robotics

Adam Stanton [1], Jared M. Moore [2]

Corresponding: Adam Stanton (a.j.stanton@aston.ac.uk)

1. School of Informatics and Digital Engineering, Aston University, Birmingham, B7 4ET, UK
2. School of Computing and Information Systems, Grand Valley State University, Allendale, MI 49401, USA

Abstract. In Evolutionary Robotics, Lexicase selection has proven effective when a single task is broken down into many individual parameterisations. Evolved individuals have generalized across unique configurations of an overarching task. Here, we investigate the ability of Lexicase selection to generalize across multiple tasks, with each task again broken down into many instances. There are three objectives: to determine the feasibility of introducing additional tasks to the existing platform; to investigate any consequential effects of introducing these additional tasks during evolutionary adaptation; and to explore whether the schedule of presentation of the additional tasks over evolutionary time affects the final outcome. To address these aims we use a quadruped animat controlled by a feed-forward neural network with joint-angle, bearing-to-target and spontaneous sinusoidal inputs. Weights in this network are trained using evolution with Lexicase-based parent selection. Simultaneous adaptation in a wall crossing task (labelled *wall-cross*) is explored when one of two different alternative tasks is also present: *turn-and-seek* or *cargo-carry*. Each task is parameterised into 100 distinct variants, and these variants are used as environments for evaluation and selection with Lexicase. We use performance in a single-task *wall-cross* environment as a baseline against which to examine the multi-task configurations. In addition, the objective sampling strategy (the manner in which tasks are presented over evolutionary time) is varied, and so data for treatments implementing uniform sampling, even sampling, or degrees of generational sampling are also presented. The Lexicase mechanism successfully integrates evolution of both *turn-and-seek* and *cargo-carry* with *wall-cross*, though there is a performance penalty compared to single task evolution. The size of the penalty depends on the similarity of the tasks. Complementary tasks (*wall-cross/turn-and-seek*) show better performance than antagonistic tasks (*wall-cross/cargo-carry*). In complementary tasks performance is not affected by the sampling strategy. Where tasks are antagonistic, uniform and even sampling strategies yield significantly better performance than generational sampling. In all cases the generational sampling requires more evaluations and consequently more computational resources. The results indicate that Lexicase is a viable mechanism for multi-task evolution of animat neurocontrollers, though the degree of interference between tasks is a key consideration. The results also support the conclusion that the naive, uniform random sampling strategy is the best choice when considering final task performance, simplicity of implementation, and computational efficiency.

Keywords: multi-objective, many-objective, evolutionary robotics, lexicase selection, transfer learning

1 Introduction

1.1 Motivation

Evolutionary Robotics (ER) uses principles of evolutionary computation to discover behaviours in artificial autonomous systems through continual adaptation of morphologies and controllers (Floreano et al., 2008). The ambitions of ER are aligned with autonomous robotics more generally, aiming to find control architectures for robots that embody a general capability to deal with problems in their worlds (Vargas et al., 2014). In the present work, we consider this general capability to mean two things: that a controller can perform well across variations of a single task ("semi-generalised control"), and that the controller is competent in multiple task domains which do not necessarily overlap. Current progress in ER means that finding controllers with competency both where multiple tasks exist and also where each task entails a multiplicity of individual parameterisations, is a desirable and feasible research objective.

1.2 Lexicase Selection

Alongside fitness metrics, genetic encodings, and other crucial components of evolutionary algorithms, the method of choosing parents for new generations, the "selection operator", is a key consideration. Early research in evolutionary computing used a single measure of fitness to select parents and the limitations of this approach ultimately highlighted the need to use selection operators that simultaneously consider a number of different dimensions of performance (Mitchell, 1998). Various evolutionary techniques exist to map and explore effective solutions in a multi-objective optimisation problem. Basic normalisation and averaging of a fitness vector producing a single scalar quantity can be applied, though this technique does not generally scale well, tending to discard high-performing solutions where they are weak in another dimension. More advanced algorithms like NSGA-II (Deb et al., 2002) explicitly acknowledge the Pareto front and archive and sort Pareto-dominant solutions. These algorithms select from the optimal front at each iteration, though they do not perform well in many-objective problems (Seada & Deb, 2015), colloquially defined as problems with more than three objectives. Lexicase selection is a novel many-objective selection operator that often selects specialists (Helmuth et al., 2020), i.e. individuals that are effective in a subset of the objective space but not necessarily the best in every objective, and is capable of scaling to at least 200 individual objectives (Moore & Stanton, 2021).

1.3 Our ER problem domain

In previous work (Moore & Stanton, 2017) we explored Lexicase selection in ER where feed-forward neural controllers are optimised to discover quadrupedal walking gaits. Animats must move towards a target while climbing over a wall of varying height. The wall is positioned half way between the animat's starting point and its target location. We call this task *wall-cross*. Variation in this obstacle constitutes the task parameterisation (i.e. the semi-generalised control problem). The obstacle can be one of 100 different heights when presented to animats, resulting in a many-objective optimization problem. Desirable gaits are those that successfully negotiate many different heights[1]. We observed that Lexicase evolved populations towards high fitness, exploring difficult areas of the solution space whilst maintaining competency on parts of the problem that had already been solved. We compared Lexicase to algorithms that were designed specifically for this problem–those presented in Stanton and Channon (2013)–and found that Lexicase, even without specific parameter tuning, outperformed those algorithms in all cases.

1.4 Expanding the domain to multiple tasks

Species in the natural world have evolved over millions of years to perform well on a variety of arbitrary problems posed by their environments: natural organisms do not evolve in response to a single, clear adaptive pressure in isolation. It is likely that these overlapping pressures are an important driver of the impressive general competences present in nature. Environmental challenges can reinforce each other and have the potential to select strongly for common adaptations that form the basic building blocks for more advanced adaptive responses. Combinations of these adaptations confer specific capabilities and ultimately respond to particular existential struggles encountered by evolving species.

With this in mind, the motivation of the present paper is to explore the performance of Lexicase selection in the ER environment outlined above while expanding the range of objectives to include a second task. Alongside the wall crossing task we introduce and explore two new problems: *turn-and-seek*, and *cargo-carry*.

In *turn-and-seek*, the wall is removed and the target position is parameterised into one of 100 configurations. Each configuration translates to a placement of the target on a 180 degree arc centered in front of the animat with performance again measured by the proximity to this target.

Figure 1 depicts an animat in each of the three task environments. In the *cargo-carry* task, animats are instantiated with a weighted box on their torso. Successful animats in this task have gaits that maintain a relatively stable torso and thus are able to carry the cargo without it falling to the ground. Here, performance is measured by time. Animats accrue fitness until the box falls from the torso to the floor, at which point the simulation is terminated. The task is parameterised by cargo mass: each of the 100 environments of this task simulates a different weight for the animat to carry.

[1]It is important to note that in this configuration, animats do not have information about the wall height. Their gaits are blind and driven only by proprioceptive, directional, vestibular and spontaneous cyclical input signals.

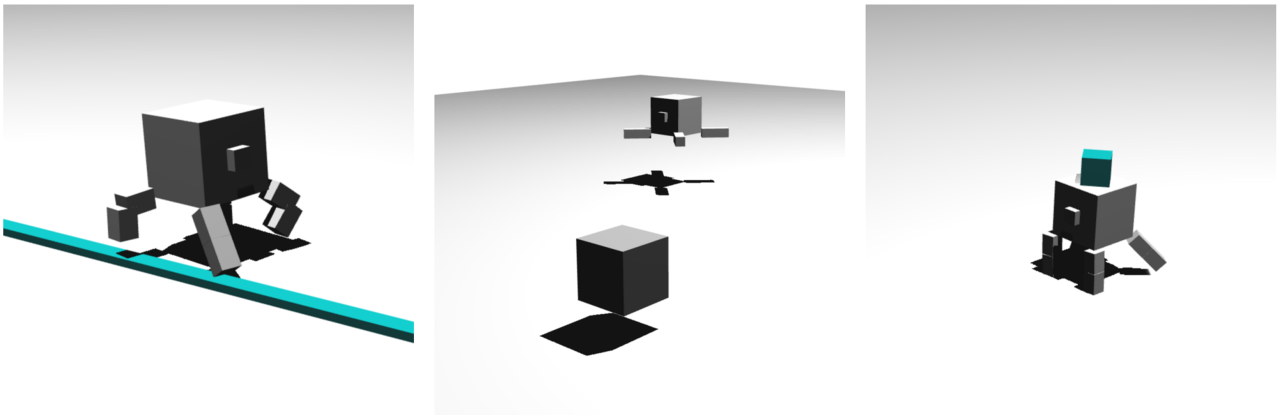

Figure 1: Examples of animat evaluation in the three task environments: *wall-cross* (left); *turn-and-seek* (centre); *cargo-carry* (right). Wall-cross and turn-and-seek both have a target that the animat is navigating toward. This can be seen in the turn-and-seek task as the cube in the foreground.

1.5 Objectives

In the present study there are three main research objectives. First, we wish to determine the feasibility of introducing the additional tasks outlined above to the existing platform. This is in terms of the practicality of integrating them into the algorithm and also to highlight any problems that can occur with the approach.

The second aim is to investigate any consequential effects of introducing these tasks during evolutionary adaptation. These consequences could include reduced performance in one or more tasks or, conceivably, positive synergistic effects if species discover broad, underlying competencies supporting multiple tasks which are threshold discoveries that open domains of even higher fitness.

Third, we aim to explore whether the schedule of presentation of the additional tasks over evolutionary time affects the final outcome. We consider whether a) structuring the presentation of tasks on a generation-by-generation basis, or b) enforcing a certain amount of time spent on each task, has effects on the final outcome of the algorithm in terms of performance of evolved species in specific tasks as well as to the overall performance across the complete problem space.

Broadly, we aim to present a first attempt to show the utility of Lexicase selection in multi-task ER problems. We provide a comparative analysis, contrasting the outcomes of evolution in the multi-task environments with those of single-task populations, a discussion of these results in terms of the interactions between the tasks and subtasks involved, and suggestions for future research directions and expansion of these ideas.

2 Background and Related Work

Evolved robot controllers have proven effective in legged locomotion (Baydin, 2012; Clune et al., 2009; Nolfi & Floreano, 2000) including transferring evolved controllers to physical systems (Koos et al., 2010; Ruud et al., 2016) with fitness often based on the distance traveled in a fixed amount of time. Increasingly, secondary considerations like damage mitigation, and generalizability of behaviours (Pinville et al., 2011) encourage the use of multi-objective algorithms considering multiple performance metrics. Biological observations further enhance systems by bringing in objectives related to efficiency of locomotion (Moore & McKinley, 2016). Subsumption architectures (Brooks, 1986; Koza, 1994; Lessin et al., 2013), behavioural diversity approaches (Doncieux & Mouret, 2013) and the combinatorial multi-objective evolutionary algorithm (Huizinga & Clune, 2021) have demonstrated controllers capable of multiple behaviours in one platform. Generalizing controllers spans learning and reacting to environmental contexts across many environments (Lehman et al., 2013), adapting and reconfiguring morphology in response to damage (Kriegman et al., 2019), and exhibiting multiple gaits for one morphology (Cully et al., 2015). Evolving *distinct* behaviours in one controller remains a long-standing goal in ER.

Adding generalizability to evolved controllers typically involves moving towards multiple fitness metrics. Multi-objective and many-objective algorithms like Lexicase selection enable scaling fitness objectives into the tens or hundreds of individual objectives. In this study, we expand on earlier investigations (Moore & Stanton, 2017, 2018, 2019, 2020, 2021) by adding new meta-tasks in addition to wall crossing, evaluating the performance of evolved individuals and investigating the performance of Lexicase selection. Adding a second task could lead the evolutionary process to new areas in the search space resulting in higher performance (Wagner et al., 2020). Switching between environments can also lead to more effective overall performance across tasks (Canino-Koning et al., 2019; Nahum et al., 2017). Understanding the underlying mechanisms that drive Lexicase's performance, and especially different parameterizations (Hernandez et al., 2022; La Cava et al., 2016) is critical to applying Lexicase effectively. The large search space of many-objective problems can create a computational challenge. Downsampling the number of objectives for consideration during Lexicase selection (Helmuth & Spector, 2020; Hernandez et al., 2019) reduces computational overhead by limiting a selection event to a subset of the objective space. We downsample to 10 objectives from a possible 200, consistent with previous wall crossing experiments.

3 Methods

Parameters for the animat, controller, two of three environments, and evolutionary algorithm have been maintained from earlier work (Moore & Stanton, 2021). The software used for these experiments is publicly available and linked at the end of this paper.

Animat Morphology The quadrupedal animat, shown in Figure 1, has a cube-shaped torso with legs placed at the four lower corners. Each leg has a 2-degree of freedom (DOF) hip and 1-DOF knee. Hips move laterally and vertically allowing the leg to go from straight out from the torso to completely vertical. Knees allow the legs to curl under and towards the torso. Animats get feedback on their position relative to a target through two sensors placed on either side of the torso.

Animat Controller A feed-forward artificial neural network (ANN) provides control signals for the joints consistent with prior investigations (Moore & Stanton, 2020, 2021). 16 inputs comprising 2 periodic oscillating signals, 2 position sensor signals, and 12 for feedback from the joints provide information about the current state of the animat within the simulation. ANNs have one 12 node hidden layer and 12 outputs to control each joint. Genomes are encoded as 336 evolvable weights.

Task Overview Three primary tasks comprise the objectives in this paper: *wall-cross*, *turn-and-seek*, and *cargo-carry*. Each of the primary tasks is subdivided into 100 unique instances resulting in a total of 300 possible objectives. In *wall-cross*, animats are evaluated on their ability to navigate to a target placed in front of the animat beyond a wall. Wall height ranges over 100 values from almost non-existent to the height of the animat's hip, see Figure 2 (left). Figure 2 shows *turn-and-seek*, individuals navigate to a target placed on an arc from the animat's left to right depending on the specific instance. *cargo-carry* places a box of varying density on top of the animat where the box must be carried for as long as possible. Due to the nature of the other two tasks, animats will still attempt to walk toward a target but fitness is not scored based on distance. Rather, fitness is the total time the box is kept aloft, with a possible maximum value of 20 as simulations are conducted for 20 seconds.

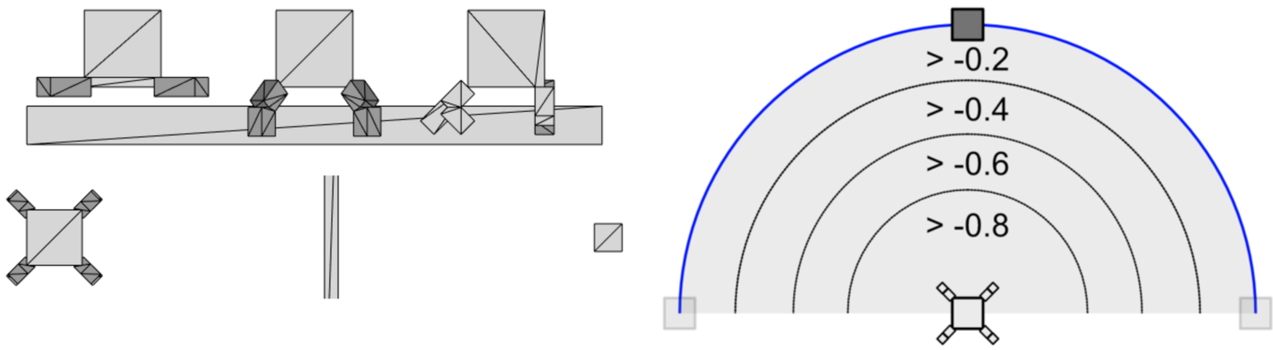

Figure 2: (Left) The maximum wall height in *wall-cross* is set at the animat's hip height if the leg is perpendicular to the ground. Initial position of the animat, wall, and target are shown from a top-down view. (Right) *turn-and-seek* is divided into 100 sub-objectives with the target placed on a semi-circle at $1.8°$ gradations. Fitness is how close the animat is to the target at a fixed timepoint. (Not to scale.) Radial lines indicate how fitness increases as the animat moves toward the target. Fitness is ultimately a straightline distance to target depending on the specific turn-and-seek environment being evaluated. Figures originally presented in (left) Moore and Stanton, 2020 and (right) Moore and Stanton, 2021.

The three tasks together are likely to facilitate some transfer of behaviors as each requires locomotion to be effective. *wall-cross* and *turn-and-seek* are complimentary as effective locomotion and the ability to navigate to a target are behaviors that lead to high performance. However, *cargo-carry* has the potential to be antagonistic to the other two tasks as stable locomotion is favored to elicit high performance in keeping the cargo on top of the animat. This is in opposition to *wall-cross* and *turn-and-seek* which have an emphasis on fast locomotion to reach the target within the simulation time. Prior investigation by Dolson et al., 2018 suggests that Lexicase selection may be sensitive to antagonistic objectives.

Evolutionary Algorithm Downsampled ϵ-Lexicase Selection (Helmuth & Spector, 2020; La Cava et al., 2016) is the evolutionary algorithm used in this study. Individual treatments comprise 20 replicate runs, each with a unique starting seed. 5,000 generations of evolution are run on populations of 50 individuals. A selection event consists of a subsample of five individuals compared on up to 10 objectives from the treatment's objective space. An ϵ- of 10% is applied to comparisons between individuals. As long as an individual is above 90% of the performance of the best individual in the sample for the objective under consideration, it is considered to be tied with the best individual and will move on to performance evaluation on the next objective. If two or more individuals still remain under consideration after all 10 downsampled objectives have been evaluated, a random selection of the remaining individuals in the subsample occurs and a tie event is recorded.

Objective Sampling Strategies In this paper, we are interested in how multiple tasks, and the interaction between them on evolved controllers, impact performance of Lexicase selection. Treatments are structured based on the tasks they include and how they sample specific task instances in the objective space to best elicit generalized controllers. Three sampling strategies alter how objectives are selected during Lexicase. Note that numbers in the sampling strategy indicate how many tasks are included. The sampling strategies are as follows:

1. *naive_2t* is the baseline, sampling 10 objectives uniformly per generation across the tasks included in a treatment. Objectives are randomly shuffled so that tasks do not always appear in the same order during Lexicase which would bias the process. Naive sampling is the default behaviour in Lexicase selection.

2. *even-shuf_2t* samples 5 objectives from each of two tasks per generation. Objectives are then shuffled randomizing ordering during selection preventing one task always appearing first.

3. *flipN_2t* sample 10 objectives per generation from only one of the two tasks. Once N generations have been selected from this task, a "flip" occurs after which objectives are sampled from the second task resulting in a changing evaluation environment during evolution. For example, with $N = 50$, objectives from the first task are selected for 50 generations before the next 50 generations are sampled from the second task. Flipping at every generation (*flip1_2t*) and every 50 generations (*flip50_2t*) are investigated in this study.

Table 1 summarises the treatments explored in this study which are variations of the objective sampling strategies discussed above.

Label	Sampling Strategy	Task 1	Task 2
naive_1t	Naive	wall-cross	NA
naive_2t_wt	Naive	wall-cross	turn-and-seek
naive_2t_wt-10000	Naive, 10k gens	wall-cross	turn-and-seek
naive_2t_wc	Naive	wall-cross	cargo-carry
even-shuf_2t_wt	Even Shuffle	wall-cross	turn-and-seek
even-shuf_2t_wc	Even Shuffle	wall-cross	cargo-carry
flip1_2t_wt	Flip every generation	wall-cross	turn-and-seek
flip1_2t_wc	Flip every generation	wall-cross	cargo-carry
flip50_2t_wt	Flip every 50 gens	wall-cross	turn-and-seek
flip50_2t_wc	Flip every 50 gens	wall-cross	cargo-carry

Table 1: Summary of treatments in this study. 20 replicates are undertaken for each treatment. *wt* indicates *wall-turn* environments while *wc* indicates *wall-cargo* environments.

4 Results

4.1 Wall Crossing Performance One Task vs Two Task

In (Moore & Stanton, 2018, 2019, 2020) we investigated characteristics of Lexicase selection in quadrupedal animats evolved for the single task of wall crossing, across 100 unique wall heights (objectives). *naive_1t* provides a baseline of performance replicating results from earlier investigations. Effective wall crossing in these individuals involves both crossing an obstacle of varying height while navigating to the target placed on the opposite side of the wall. Evolved individuals thus have some ability to navigate and locomote which is presumably beneficial in the turning task and also likely detrimental in the cargo task since movement can cause the cargo to fall to the ground. To establish benchmarks of how individuals only evolved for wall crossing versus those evolved to perform wall crossing and a second task, we investigate combinations of two environments as follows.

Figure 3 plots the performance of the best individual per replicate for *naive_1t* and *naive_2t_wt* treatments in wall crossing. (The “1t” in *naive_1t* indicates evolution with only one task while “2t” signifies two tasks; “wt” in *naive_2t_wt* represents wall crossing and turn-and-seek as the two evolutionary tasks.) Blue and red dots indicate outliers from the respective treatment’s boxplot distribution both positive or negative. Each box represents the distribution of the best mean performance individual per replicate on the specific objective. Performance for the lower wall heights is nearly identical with most replicates able to reach the target. As wall height increases towards middle heights, individuals evolved only for wall crossing evolve higher performance than those evolved for both wall crossing and turning. Wall crossing ability tapers off in both as wall height reaches the upper limits, wherein very specific gaits must evolve to cross these challenging obstacles. *naive_1t* does significantly outperform *naive_2t_wt* in the wall crossing task. Statistical significance is determined by a Wilcoxon rank-sum test with Bonferroni correction throughout this paper.

Figure 4 plots the performance of the best individual per replicate in the turning task for *naive_1t* and

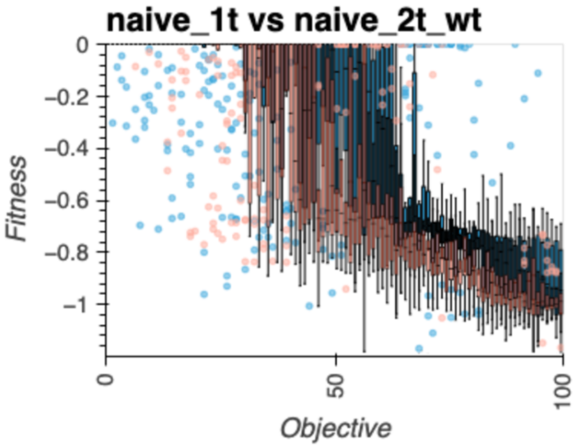

Figure 3: *naive_1t* (blue) and *naive_2t_wt* (red) best individual by mean performance per replicate in wall crossing task. Fitnesses below -0.4 indicate individuals that are unable to cross the wall. Figure adapted from Moore and Stanton, 2021.

naive_2t_wt. Although not exposed to the turning task during evolution, many of the individuals evolved only for wall crossing in *naive_1t* are able to navigate to the target when it is placed nearly in front of the animat (objectives 40-60). As the target is placed further to the left or right, performance tapers off as these target locations are quite different than those encountered during evolution in wall crossing. *naive_2t_wt* significantly outperforms *naive_1t* with performance reaching nearly optimal by generation 500. Still, the ability to navigate to targets placed near the center for individuals in *naive_1t* suggest that the tasks of wall crossing and turn-and-seek may be complementary. That is, behaviours evolved for one task may assist in solving the other.

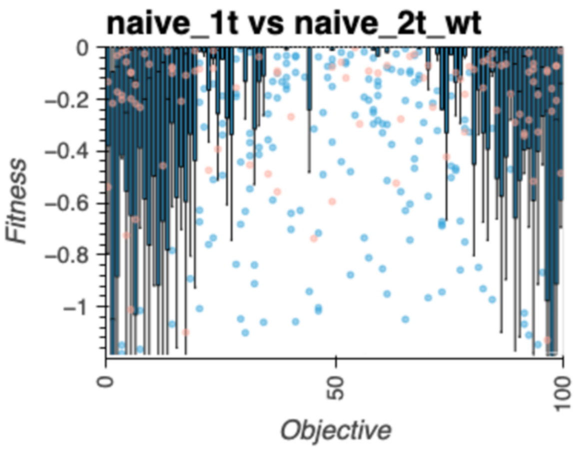

Figure 4: *naive_1t* (blue) and *naive_2t_wt* (red) best individual by mean performance per replicate in turning task. *naive_2t_wt* evolves near perfect turning performance with only a few outliers not evolving full generalization on this task. Figure adapted from Moore and Stanton, 2021.

naive_1t and *naive_2t_wt* are both evolved for 5,000 generations even though *naive_2t_wt* has 200 objectives across two tasks. The increase in the number of unique objectives might mean performance differences in wall crossing are due to fewer selection events occurring for the wall crossing task. We conduct a second two task naive treatment, *naive_2t_wt-10000*, to see if evolving for 10,000 generations allows for similar performance in wall crossing. Figure 5 plots wall crossing performance for the best individual per replicate. No significant difference in wall crossing performance arises between *naive_1t* and *naive_2t_wt-10000*. Performance between *naive_2t_wt* and *naive_2t_wt-10000* is also not significantly different. Naive selection strategies are able to evolve similar wall crossing performance, while also incorporating a second task when allowed to evolve for similar generations relative to the number of tasks. However, our goal with the other sampling strategies is to assess whether it is possible to evolve generalized two task performance in the same number of generations as one task performance. Thus, we only evolve for 5,000 generations in the remaining treatments.

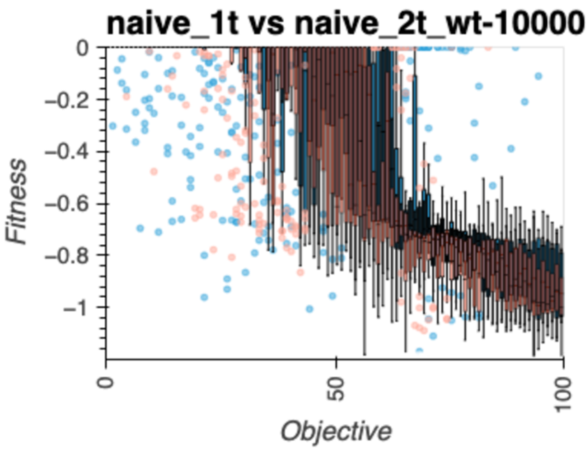

Figure 5: *naive_1t* (blue) and *naive_2t_wt-10000* (red) best mean performance individual per replicate in wall crossing. No significant difference in performance between treatments. Figure originally presented in Moore and Stanton, 2021.

Cargo carrying provides another task to evolve individuals alongside wall crossing. Figure 6 plots the wall crossing performance for individuals evolved only for wall crossing in *naive_1t* and those evolved for wall crossing and cargo in *naive_2t_wc* ("wc" in *naive_2t_wc* represents wall crossing and cargo carrying as the two tasks against which agents are evaluated.) Performance in wall crossing of the best individuals evolved for wall crossing and cargo is significantly lower than that of individuals evolved only for wall crossing after 5,000 generations. Furthermore, Figure 7 plots the cargo carrying performance for the best individuals in *naive_1t* and *naive_2t_wc*. Performance between the two treatments is significantly different with *naive_2t_wc* evolving near perfect performance in the cargo carrying task while *naive_1t* individuals have no ability to carry the cargo no matter the specific cargo density. It takes some time for the box to reach the ground resulting in fitnesses near 3 for *naive_1t*.

For cargo carrying, individuals are not evaluated based on their distance from a target as in wall crossing and turning. Instead, individuals are measured on their ability to prevent the cargo from falling to the ground in terms of how many seconds they keep it aloft. The difference in selective pressure results in there being an antagonistic relationship between wall crossing and cargo carrying. This push-pull between the tasks invites the hypothesis that, compared to a naive mixing of tasks and objectives, a structured

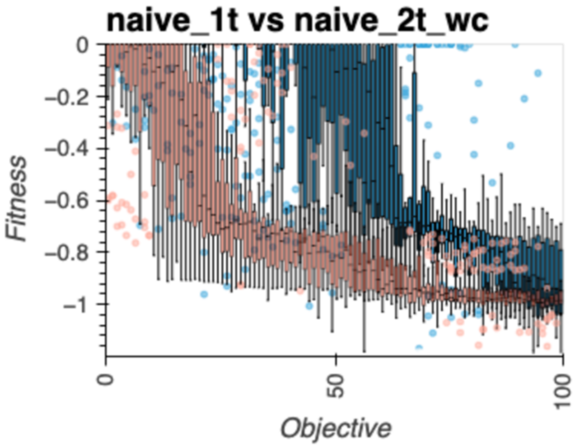

Figure 6: *naive_1t* (blue) and *naive_2t_wc* (red) best mean performance individual per replicate in wall crossing. Performance between the two treatments is significantly different.

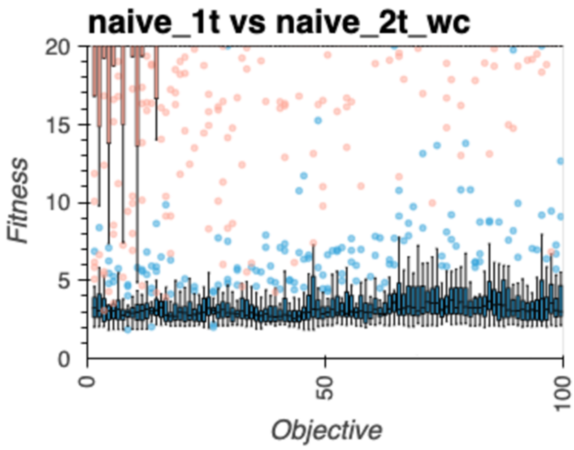

Figure 7: *naive_1t* (blue) and *naive_2t_wc* (red) best mean performance individual per replicate in cargo task. Performance between the two treatments is significantly different. Note the red bars are almost all at 20, which is maximum performance in the cargo carrying task.

presentation of objectives over evolutionary time might improve generalised performance on both tasks. The rationale is that adaptation collapses to one of the two opposing problems when tasks and objectives are well mixed; by forcing longer adaptive periods on the other task it is possible that sufficient progress is made to steer populations towards areas of the solution space with potential to achieve high performance on both problems. This could be due to a differential in the (a posteriori) difficulty of the tasks, or simply a consequence of the complex dynamics of the adaptive landscape and evolutionary process.

4.2 Two Task Complementary Treatments

Figure 8 plots the best performing individual per replicate determined by mean performance in wall crossing, turning, and mean performance for all *wt* treatments. *naive_1t* is also provided as a baseline for comparison. All two task treatments evolve turning with no significant difference in performance. Wall crossing performance is consistent for all two task treatments as well. All *wt* treatments significantly outperform *naive_1t* in the turning task.

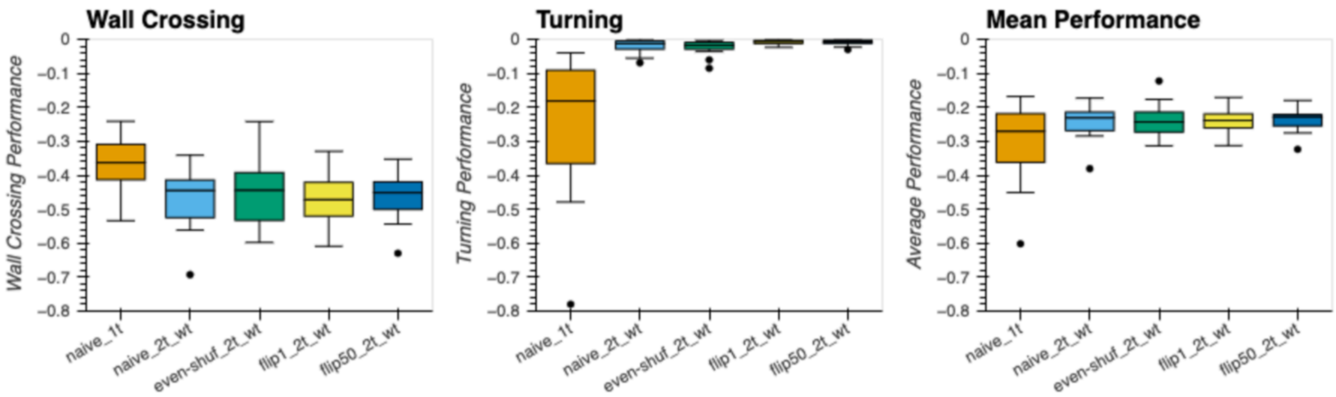

Figure 8: Best individual by mean performance across the two tasks per replicate across treatments run for 5,000 generations. Figure originally presented in Moore and Stanton, 2021.

Figure 9 plots the performance per individual objective for the best individual per replicate across the *wt* treatments. Although performance slightly lags *naive_1t* in all *wt* treatments for wall crossing, generalized behaviors are evident across *wt* treatments. All *wt* treatments evolve individuals capable of traversing low wall heights with performance tapering near the middle wall heights. This is consistent with *naive_1t* although it does have more effective performance on the upper-middle wall heights. Tall walls are difficult to traverse due to the morphology of the animat and performance tapers off accordingly for all treatments.

4.2.1 Lexicase Dynamics

Examining the dynamics of Lexicase selection across objective sampling strategies can help elicit differences that are not apparent from examining performance of evolved individuals. Figure 10 plots the total number of environments considered during selection per replicate over evolutionary time. A higher number of selection environments indicates that a sampling strategy had to go farther into the downsampled objectives to separate individuals during the Lexicase process.

Figure 11 plots the number of simulations conducted per replicate. As individuals are filtered out during selection, the total number of simulations declines. Treatments with fewer individual evaluations might be good at reducing the subsample of individuals in a selection to one or two performant individuals that are then tied during the remainder of a Lexicase selection event. Whereas, a high number of selection environments and subsequently high number of evaluations might mean a treatment is poor at filtering out individuals throughout selection. Both metrics can indicate the computational efficiency of a treatment.

Although performance among the wall and turn sampling strategies are not significantly different, both *flip* treatments exhibit considerable disparity in both number of environments considered and individual evaluations between *turn* and *wall* environments while *naive_2t_wt* and *even-shuf_2t_wt* do not. Considering that the flip treatments sample objectives from only one task per generation it appears that turning is not as effective as wall crossing at separating individuals out during the selection process.

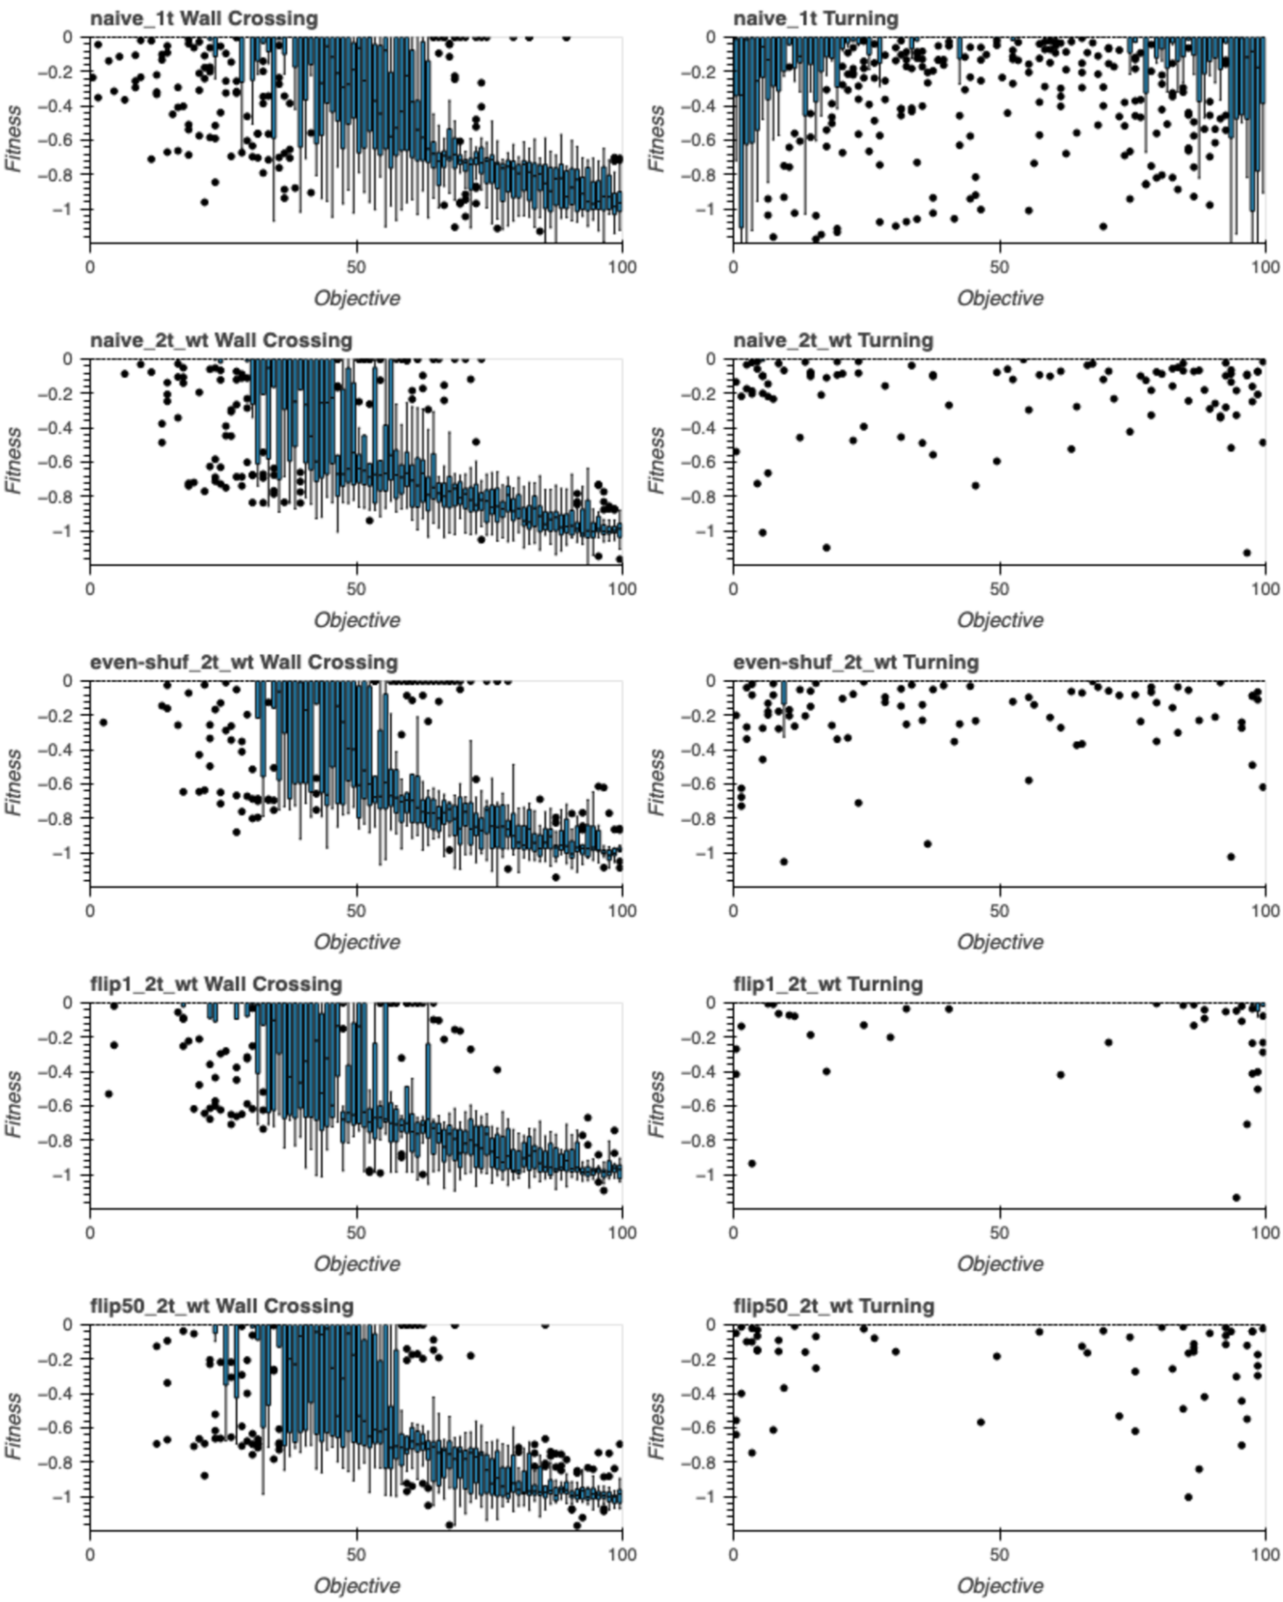

Figure 9: Performance per objective of the best individual by mean performance per replicate across *wt* treatments run for 5,000 generations.

naive_2t_wt and *even-shuf_2t_wt* do not exhibit a similar difference between the two environments. Both treatments result in a uniform sampling between environments based on the two figures even though *naive_2t_wt* does not actively enforce even sampling. Given the differences in Figures 10 and 11 for *flip* versus *non-flip* treatments, we hypothesize that the turning task is not as effective in filtering individuals as

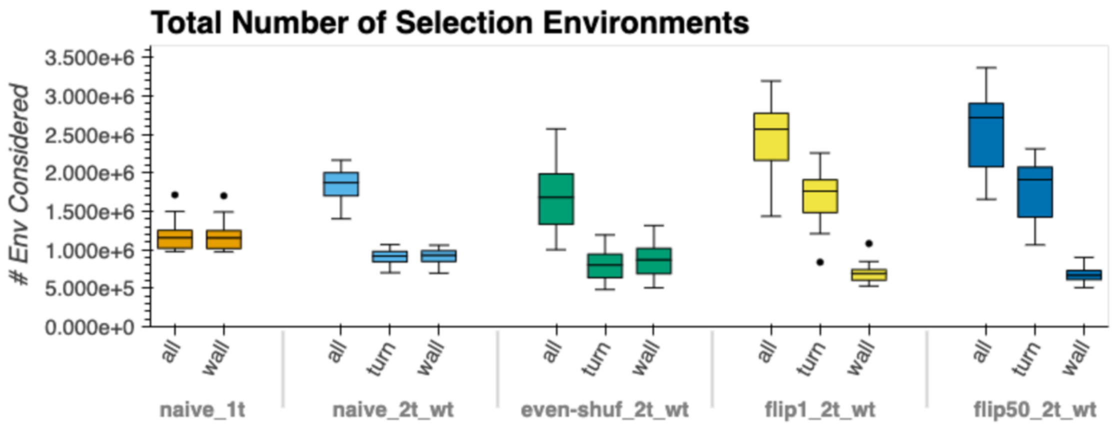

Figure 10: Count of environments considered per replicate. *all* is the sum of wall and turn. Higher numbers suggest more environments were needed per selection event during Lexicase. Figure originally presented in Moore and Stanton, 2021.

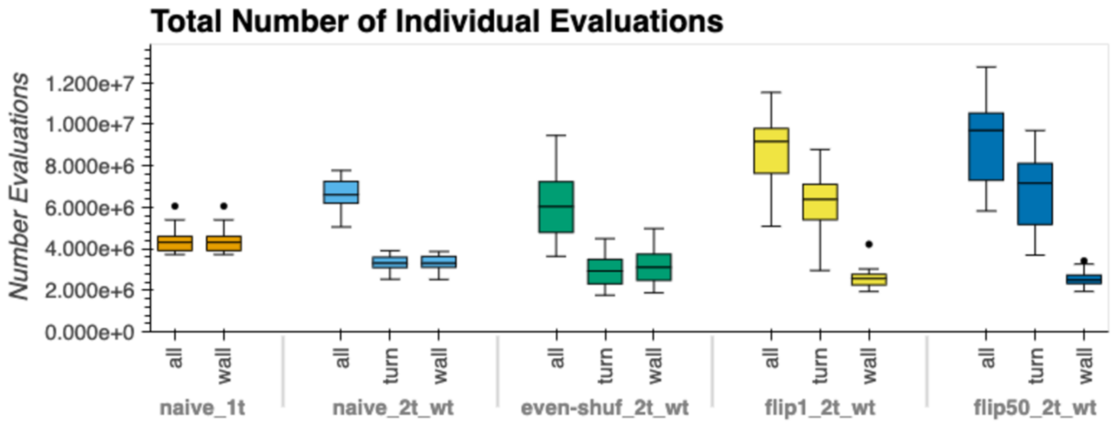

Figure 11: Count of simulations per replicate. *all* is the sum of wall and turn. Simulations represent the significant computational cost. Figure originally presented in Moore and Stanton, 2021.

the wall crossing task. Explicitly requiring that only objectives from turning be used for a generation as-in the *flip* treatments significantly increases computation time for an evolutionary run without a significant increase in performance in evolved individuals.

Figure 12 plots the number of tiebreak events across replicates. Tiebreaks represent a situation where no individual is determined as better than the others in the sampled selection subset. There is no significant difference in the number of tiebreaks between *naive_2t_wt* and *even-shuf_2t_wt*, or between *flip1_2t_wt* and *flip50_2t_wt*. Adding the turning task significantly raises the number of tiebreaks between *naive_1t* and the two-task treatments while the *flip* strategy further increases tiebreaks. This additional dynamic further suggests that *flipping* is not an efficient Lexicase selection strategy for these two tasks. Specifically, a high number of tiebreaks suggests an inability to effectively filter less performant individuals if there is not a corresponding increase in performance relative to other treatments.

Figure 13 plots the percentage of individuals filtered out in a selection event per objective on the left hand side with the number of times the objective occurs in a replicate on the right hand side. Wall crossing objectives are numbered 0-99 while turning is 100-199. The figure shows that across treatments, middle

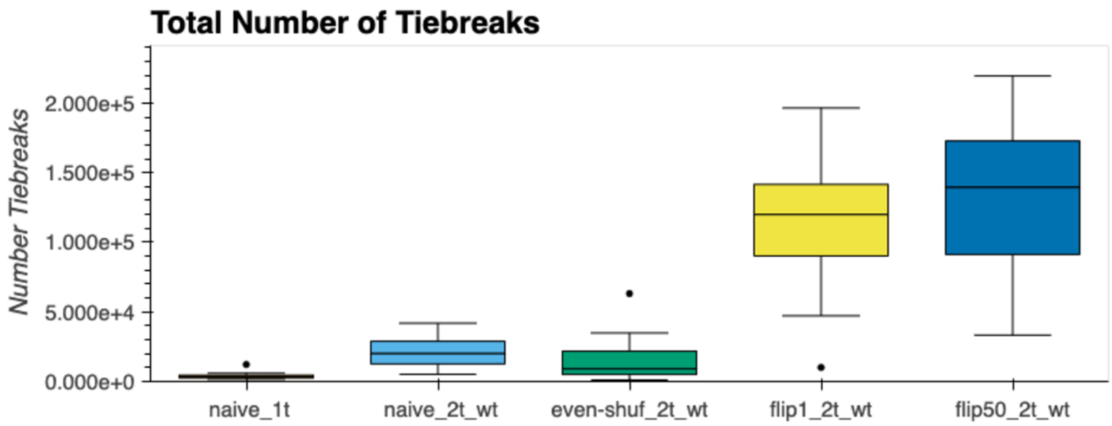

Figure 12: Count of tiebreaks across replicates. Tiebreaks arise when selection cannot isolate one individual during the Lexicase operation. Figure originally presented in Moore and Stanton, 2021.

wall heights are the most effective at reducing the number of individuals under consideration by Lexicase selection. We hypothesize that short walls are relatively easy to cross, assuming an effective locomotion strategy has evolved, while higher wall heights are also weak selectors as most individuals are incapable of crossing them. Middle wall heights reach a balance with only the most effective individuals able to cross.

Turning suggests that the far left and far right objectives are the most effective filters while targets placed in front of the quadrupeds are less effective. Targets placed in front of the quadrupeds share similar sensor feedback as those of all wall crossing environments likely resulting in some behavioural transfer between the two tasks. These front facing targets are similar to low wall height objectives in wall crossing which primarily require an effective locomotion strategy and support the low filtering effectiveness of short wall heights.

4.3 Two Task Antagonistic Treatments

Figure 14 plots the best performing individual per replicate determined by mean performance in wall crossing, cargo carrying, and mean performance for all *wc* treatments. *naive_1t* is again provided as a baseline for comparison. All two task treatments evolve cargo carrying ability with no significant difference in performance, significantly outperforming *naive_1t*. Unlike the turning task, the best evolved wall crossing individuals from *naive_1t* are not effective at carrying cargo. This suggests that transfer of evolved behaviours is reduced between *wc* when compared to results for *wt* in Figure 8. Individuals in *naive_2t_wc*, *even-shuf_2t_wc*, *flip1_2t_wc*, and *flip50_2t_wc* fail to evolve effective generalized wall crossing behaviour when compared to their associated treatments in *wt*.

Figure 15 plots the performance per individual objective for the best individual per replicate across the *wc* treatments. In contrast to *wt* treatments, performance is noticeably lower in most wall heights. *naive_2t_wc* and *even-shuf_2t_wc* are able to evolve some competency on low wall heights while the *flip* treatments show lower performance. *flip50_2t_wc* is unable to cross any walls while evolving high performance in turning.

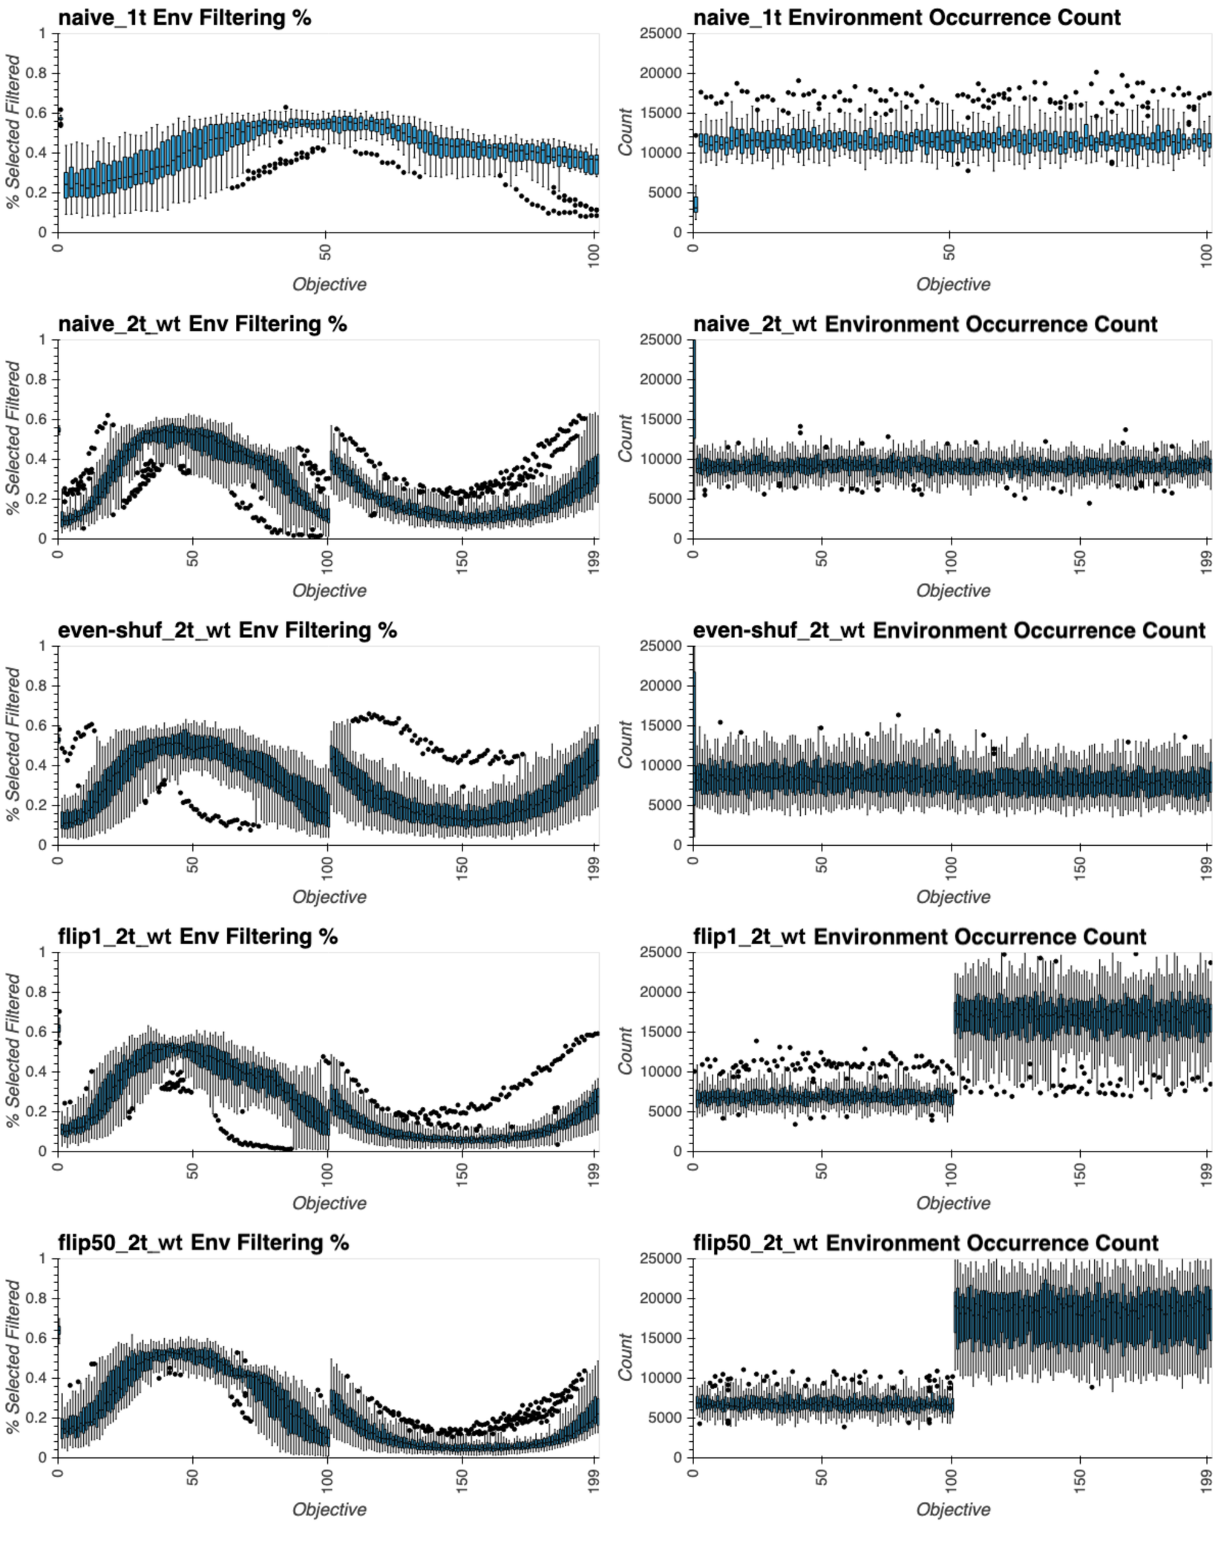

Figure 13: (Left) Filtering percentage during selection distribution across replicates. (Right) Each environment's occurrence count in selection across replicates. Figure originally presented in Moore and Stanton, 2021.

4.3.1 Lexicase Dynamics

Figures 16 and 17 plot the total number of selection environments and number of individual evaluations per replicate across treatment for the *wall* and *cargo* tasks. Although the relationship between the two tasks is different for *wc*, Lexicase dynamics are similar to *wt*. Figure 18 plots the number of tiebreak events per replicate across treatments. As in the *wt* experiment, the *flip* strategy has significantly higher number of

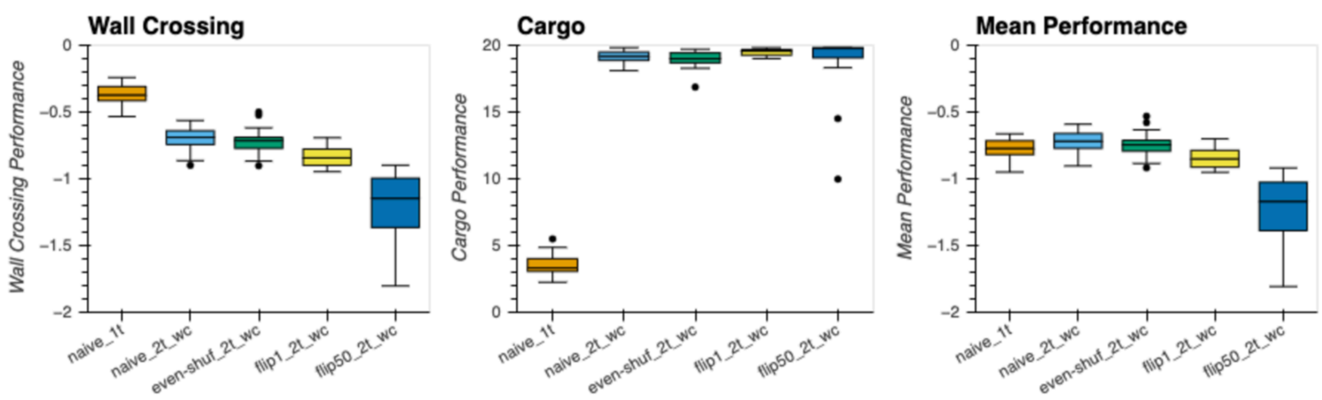

Figure 14: Best individual by mean performance across the two antagonistic tasks per replicate across treatments run for 5,000 generations. Note y-axis for wall crossing is extended to -2 to include *flip50_2t_wc* range.

tiebreaks than *naive_2t_wc* and *even-shuf_2t_wc* for two environment treatments.

Figure 19 plots (left) the percentage of individuals filtered in a selection event per objective with the (right) number of times the objective occurs in a replicate. Here, dynamics are different than the *wt* experiments as the curve of filtering effectiveness for wall crossing objectives shifts towards lower wall heights. This is especially apparent in the *flip1_2t_wc* and *flip50_2t_wc* treatments. Cargo carrying filtering effectiveness is also quite low. The high occurrence count of cargo environments in *flip* treatments shown on the right column of Figure 19, the significantly higher count of tiebreaks versus *naive* and *even shuffling*, and the high performance in the cargo task together indicate that the general solution to the *cargo-carry* task across unique box densities is relatively easy to find without behavioural transfer from *wall-cross*.

5 Discussion

Generalized behaviour is a long-standing goal in robotics. Here, we evaluate Lexicase selection on evolving quadrupedal animats in two two-task generalization problems. Four objective sampling strategies exhibit moderate losses in performance in wall crossing compared to the baseline for the *wt* problem with a larger decrease in the *wc* problem for *flipping* strategies. The small performance differences between the sampling strategies, increases in computational effort for *flip* treatments, and no need to configure additional parameters suggests that a *naive* objective sampling strategy is effective across both two task problems examined in this study.

Our first objective in this study is to assess the feasibility of integrating multiple environments into the evolutionary process with Lexicase selection. Notably, individuals in *wt* environments demonstrate both effective wall crossing on the majority of objectives and the ability to turn and seek a target across objectives. There are synergies between the two tasks that have likely been exploited in evolved individuals as we see some ability to execute turn-and-seek behaviours in *naive_1t* which does not encounter this task during

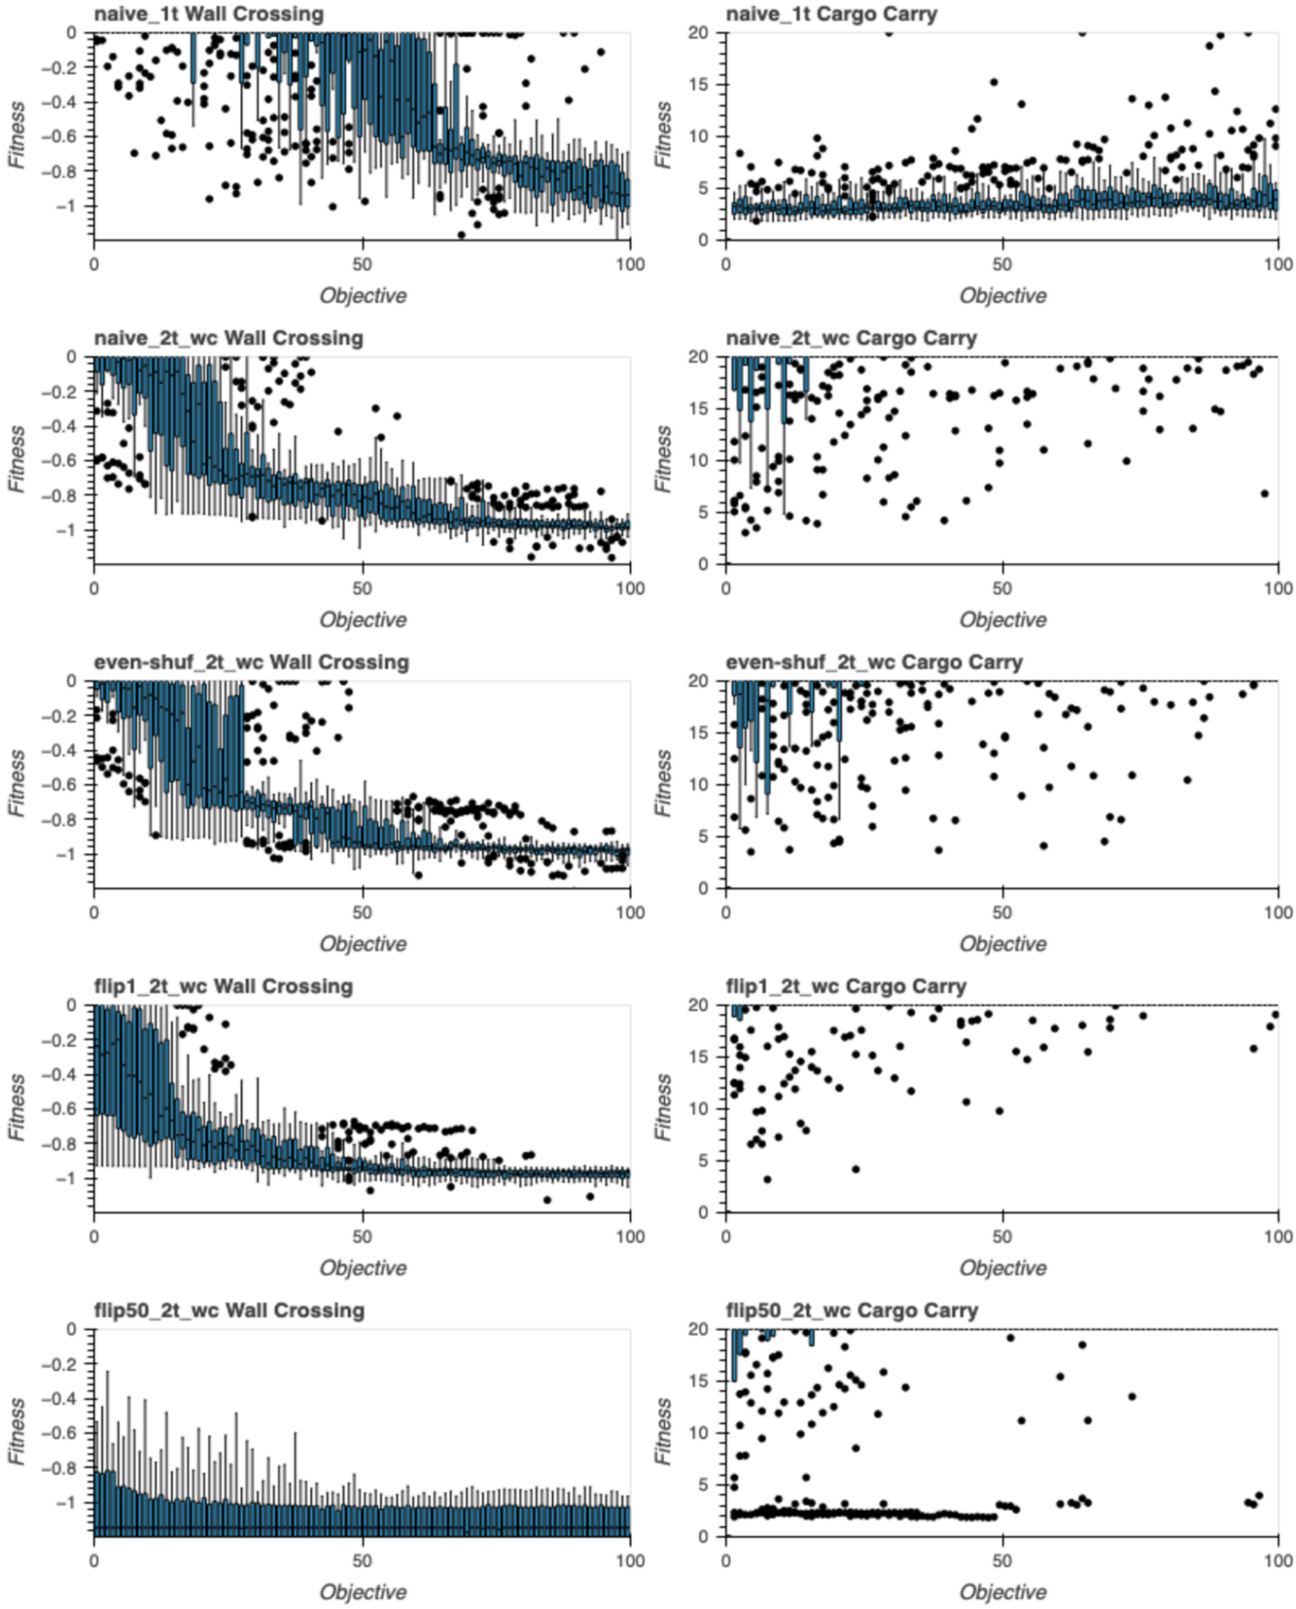

Figure 15: Performance per objective of the best individual by mean performance per replicate across *wc* treatments run for 5,000 generations.

evolution. Since both tasks have a target, evolved individuals have their ability to navigate reinforced across tasks, leading to high performance in both. Individuals evolved in the *wc* environments evolve effective wall crossing behaviours in lower wall heights only, while they are effective broadly at the cargo carrying task. Observing evolved individuals in *wc*, carrying a box tends toward small leg movements kept under the

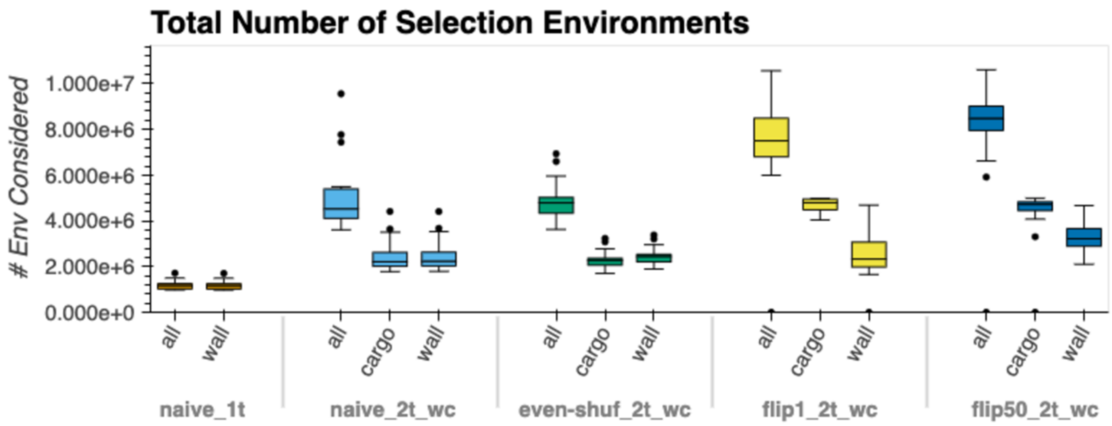

Figure 16: Count of environments considered per replicate. *all* is the sum of wall and cargo. Higher numbers suggest more environments were needed per selection event during Lexicase.

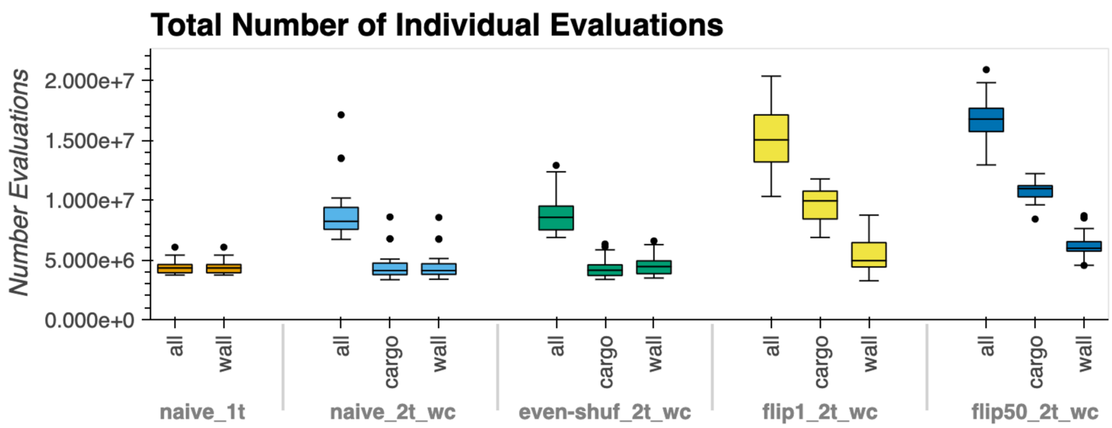

Figure 17: Count of simulations per replicate. *all* is the sum of wall and cargo. Simulations represent the significant computational cost.

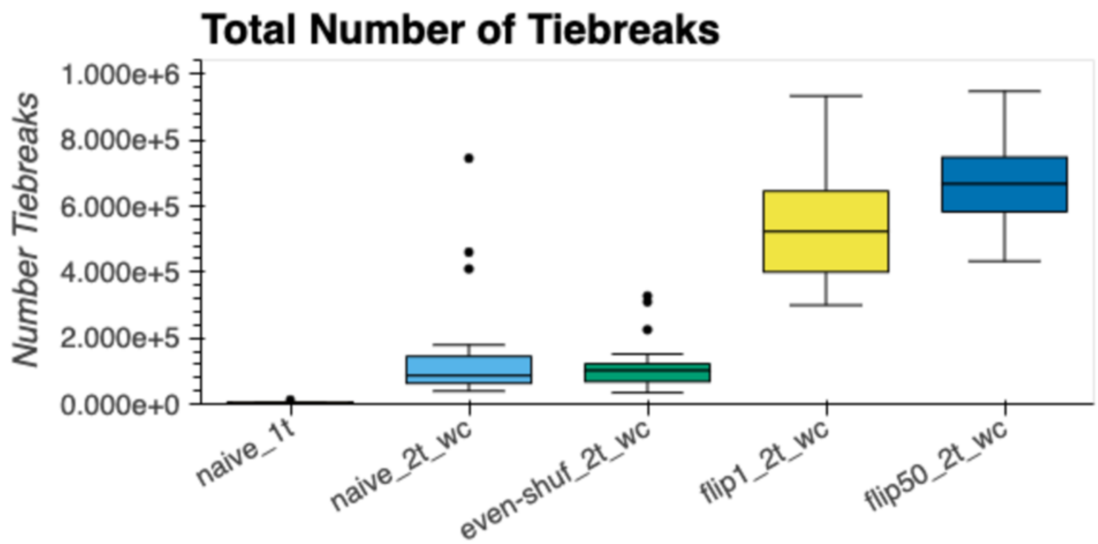

Figure 18: Count of tiebreaks across replicates for antagonistic two tasks.

torso maintaining a stable body posture whereas higher wall heights in wall crossing require sweeping leg movements that reach upwards to step over the obstacle. These conflicting evolutionary pressures likely lead to the low performance in wall crossing for *wc* treatments and the reduced generalization of individuals in *wc*

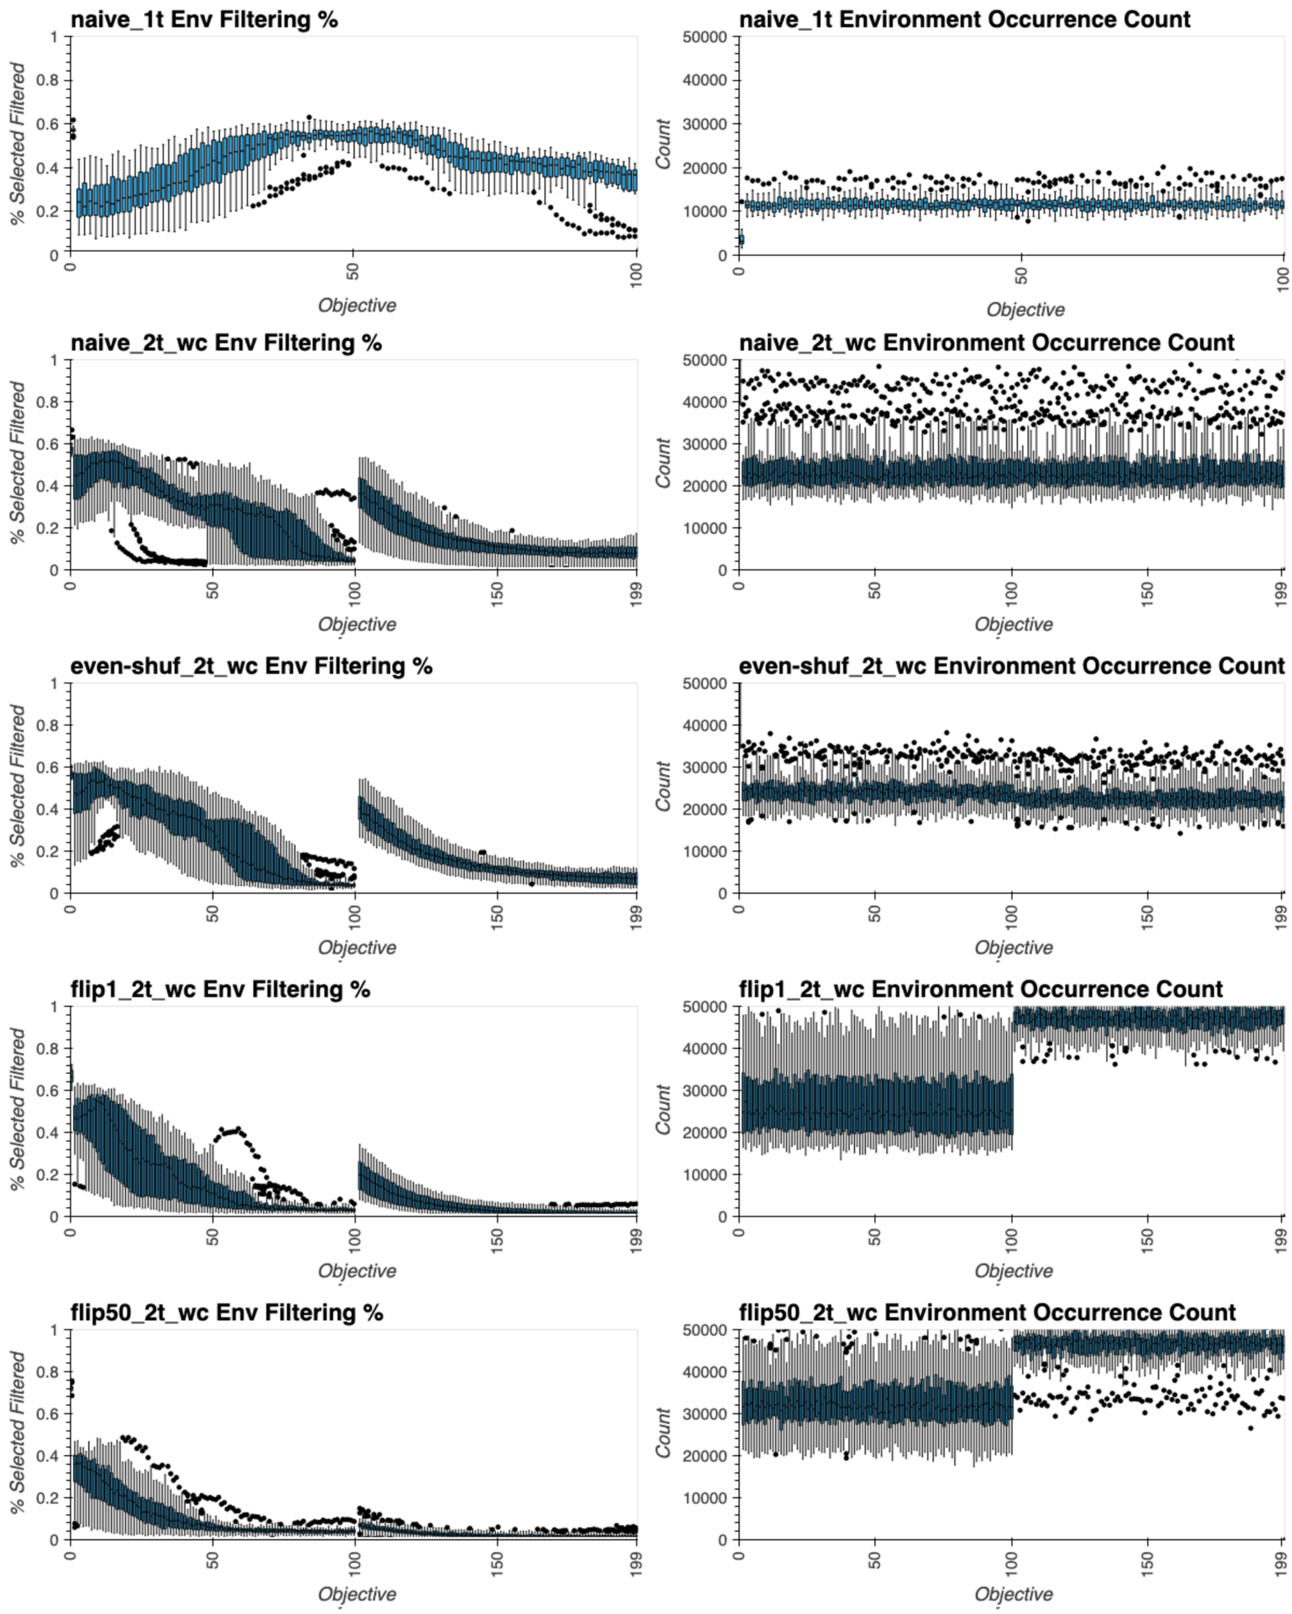

Figure 19: (Left) Filtering percentage during selection distribution across replicates. (Right) Each environment's occurrence count in selection across replicates.

versus *wt*. Still, Lexicase selection is able to integrate multiple general tasks effectively in the evolutionary process. We do observe a decrease in wall crossing performance when adding a second task in both *wt* and *wc* while keeping the number of generations of evolution stable across the one task baseline and two task treatments.

The second objective in this study is to examine any consequences of introducing the additional tasks, while the third objective explores the schedule of presentation of objectives over evolutionary time. *flipping* strategies are intended to allow the population to specialize on an environment for a specified number of generations. In *flip1_2t_wt* and *flip1_2t_wc*, objectives were sampled from one task per generation with the task flipping every generation. This was intended to prevent one task from dominating selection resulting in poor performance in one task. *flip50_2t_wt* and *flip50_2t_wc* reduced the flipping rate to alternating tasks every 50 generations. Here, individuals would be given a substantial amount of time to specialize on an objective, hypothetically increasing their competency on a task before swapping to the other task. *flipping* strategies appear to be broadly poorer than *naive* or *even shuffling* as their performance is either level with, see Figure 8, or substantially reduced across all environments, see Figure 14, considered in this study. Moreover, when tasks are antagonistic as in the case of *wc*, it appears that *flip50_2t_wc* overspecializes on the cargo task, resulting in a failure to generalize to the wall crossing task. Whereas, *flip1_2t_wc* is able to maintain slightly lower performance than *naive_2t_wc* and *even-shuf_2t_wc* in wall crossing. Against of our intuition that allowing replicates to evolve in one environment for an extended number of generations would improve competency on the tasks, the extra time given to specialize may in fact be detrimental to generalization. This result could be due to either catastrophic forgetting of the second task or evolution of a specialized behaviour for one task that preventing effective behaviours in the second task. In addition, there is a significantly higher number of evaluations and tiebreaks in both two environment combinations when compared to *naive* or *even-shuffling* strategies.

6 Future Work

In future work, we plan to expand the number of tasks while exploring how to better integrate tasks that introduce conflicting selective pressures. Of specific interest will be any changes in performance or characteristics of Lexicase selection that might arise in those broader search spaces. We also plan to leverage the filtering efficacy of objectives to explore whether an *adaptive* strategy favoring individual objectives might enhance performance of the algorithm in large search spaces.

7 Acknowledgements

The authors thank William La Cava, Thomas Helmuth, Edward Pantridge, Emily Dolson, and Lee Spector for their conversations about Lexicase selection. We would also like to acknowledge Aston University, Keele University, and Grand Valley State University for supporting this research. The software used for this project has been made available via the public repository linked below[2].

[2] https://github.com/machine-machines/lexicase-alife-journal-2022

References

Baydin, A. G. (2012). Evolution of central pattern generators for the control of a five-link planar bipedal walking mechanism. *Paladyn. Journal of Behavioral Robotics*, *3*(1), 45–53.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, *2*(1), 14–23.

Canino-Koning, R., Wiser, M. J., & Ofria, C. (2019). Fluctuating environments select for short-term phenotypic variation leading to long-term exploration. *PLOS Computational Biology*, *15*(4), 1–32. https://doi.org/10.1371/journal.pcbi.1006445

Clune, J., Beckmann, B. E., Ofria, C., & Pennock, R. T. (2009). Evolving coordinated quadruped gaits with the HyperNEAT generative encoding, In *Proceedings of the ieee congress on evolutionary computation*, Trondheim, Norway, IEEE.

Cully, A., Clune, J., Tarapore, D., & Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503–507. https://doi.org/10.1038/nature14422

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, *6*(2), 182–197. https://doi.org/10.1109/4235.996017

Dolson, E., Banzhaf, W., & Ofria, C. (2018). Applying ecological principles to genetic programming (W. Banzhaf, R. S. Olson, W. Tozier, & R. Riolo, Eds.). In W. Banzhaf, R. S. Olson, W. Tozier, & R. Riolo (Eds.), *Genetic programming theory and practice xv*, Springer International Publishing.

Doncieux, S., & Mouret, J. B. (2013). Behavioral diversity with multiple behavioral distances, In *Proceedings of the 2013 ieee congress on evolutionary computation*, Cancun, Mexico, IEEE.

Floreano, D., Husbands, P., & Nolfi, S. (2008). Evolutionary Robotics. In *Handbook of Robotics*. Berlin, Springer Verlag.

Helmuth, T., Pantridge, E., & Spector, L. (2020). On the importance of specialists for lexicase selection. *Genetic Programming and Evolvable Machines*, *21*(3), 349–373. https://doi.org/10.1007/s10710-020-09377-2

Helmuth, T., & Spector, L. (2020). Explaining and exploiting the advantages of down-sampled lexicase selection, In *The 2020 conference on artificial life: A hybrid of the european conference on artificial life (ecal) and the international conference on the synthesis and simulation of living systems (alife)*, Montreal, Quebec, Canada, MIT Press. https://doi.org/10.1162/isal_a_00334

Hernandez, J. G., Lalejini, A., Dolson, E., & Ofria, C. (2019). Random subsampling improves performance in lexicase selection, In *Proceedings of the genetic and evolutionary computation conference companion*, Prague, Czech Republic, ACM. https://doi.org/10.1145/3319619.3326900

Hernandez, J. G., Lalejini, A., & Ofria, C. (2022). An exploration of exploration: Measuring the ability of lexicase selection to find obscure pathways to optimality. In W. Banzhaf, L. Trujillo, S. Winkler, & B. Worzel (Eds.), *Genetic programming theory and practice xviii* (pp. 83–107). Singapore, Springer Singapore. https://doi.org/10.1007/978-981-16-8113-4_5

Huizinga, J., & Clune, J. (2021). Evolving Multimodal Robot Behavior via Many Stepping Stones with the Combinatorial Multi-Objective Evolutionary Algorithm. *Evolutionary Computation*, https://direct.mit.edu/evco/articl pdf/doi/10.1162/evco_a_00301/1974750/evco_a_00301.pdf, 1–34. https://doi.org/10.1162/evco_a_00301

Koos, S., Mouret, J. B., & Doncieux, S. (2010). Crossing the reality gap in evolutionary robotics by promoting transferable controllers, In *Proceedings of the 2010 acm genetic and evolutionary computation conference*, Portland, Oregon, USA, ACM. https://doi.org/10.1145/1830483.1830505

Koza, J. R. (1994). Evolution of a subsumption architecture that performs a wall following task for an autonomous mobile robot, In *Proceedings of the workshop on computational learning theory and natural learning systems : Intersections between theory and experiment*, Berkeley, California, USA, MIT Press.

Kriegman, S., Walker, S., Shah, D. S., Kramer-Bottiglio, R., & Bongard, J. C. (2019). Automated shapeshifting for function recovery in damaged robots (A. Bicchi, H. Kress-Gazit, & S. Hutchinson, Eds.). In A. Bicchi, H. Kress-Gazit, & S. Hutchinson (Eds.), *Robotics: Science and systems xv, university of freiburg, freiburg im breisgau, germany, june 22-26, 2019*, University of Freiburg. https://doi.org/10.15607/RSS.2019.XV.028

La Cava, W., Spector, L., & Danai, K. (2016). Epsilon-lexicase selection for regression, In *Proceedings of the genetic and evolutionary computation conference 2016*, Denver, Colorado, USA, ACM.

Lehman, J., Risi, S., D'Ambrosio, D., & Stanley, K. O. (2013). Encouraging reactivity to create robust machines. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, *21*(6), 484–500. https://doi.org/10.1177/1059712313487390

Lessin, D., Fussell, D., & Miikkulainen, R. (2013). Open-ended behavioral complexity for evolved virtual creatures, In *Proceedings of the 2013 acm genetic and evolutionary computing conference*, Amsterdam, Netherlands, ACM.

Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT press.

Moore, J. M., & McKinley, P. K. (2016). A comparison of multiobjective algorithms in evolving quadrupedal gaits (E. Tuci, A. Giagkos, M. Wilson, & J. Hallam, Eds.). In E. Tuci, A. Giagkos, M. Wilson, & J. Hallam (Eds.), *From animals to animats 14: Proceedings of the 14th international conference on simulation of adaptive behavior, sab 2016*, Aberystwyth, UK, Springer International Publishing.

Moore, J. M., & Stanton, A. (2017). Lexicase selection outperforms previous strategies for incremental evolution of virtual creature controllers, In *Proceedings of the 14th european conference on artificial life*, Lyon, France, MIT Press.

Moore, J. M., & Stanton, A. (2018). Tiebreaks and diversity: Isolating effects in lexicase selection, In *Proceedings of the 16th international conference on the simulation and synthesis of living systems*, Tokyo, Japan, ACM.

Moore, J. M., & Stanton, A. (2019). The limits of lexicase selection in an evolutionary robotics task, In *The 2019 conference on artificial life: A hybrid of the european conference on artificial life (ecal) and the international conference on the synthesis and simulation of living systems (alife)*, Newcastle-Upon-Tyne, UK, MIT Press.

Moore, J. M., & Stanton, A. (2020). When specialists transition to generalists: Evolutionary pressure in lexicase selection, In *Artificial life conference 2020 proceedings*, Online. MIT Press.

Moore, J. M., & Stanton, A. (2021). Objective sampling strategies for generalized locomotion behavior with lexicase selection, In *Artificial life conference 2021 proceedings*, Online, MIT Press. MIT Press. https://doi.org/10.1162/isal_a_00398

Nahum, J. R., West, J., Althouse, B. M., Zaman, L., Ofria, C., & Kerr, B. (2017). Improved adaptation in exogenously and endogenously changing environments. https://doi.org/10.1162/isal_a_052

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence and technology of self-organizing machines*. The MIT Press.

Pinville, T., Koos, S., Mouret, J.-B., & Doncieux, S. (2011). How to promote generalisation in evolutionary robotics: The progab approach, In *Proceedings of the 13th annual conference on genetic and evolutionary computation*, Dublin, Ireland, ACM. https://doi.org/10.1145/2001576.2001612

Ruud, E. L., Samuelsen, E., & Glette, K. (2016). Memetic robot control evolution and adaption to reality, In *Proceedings of the 2016 ieee symposium series on computational intelligence (ssci)*, Athens, Greece, IEEE.

Seada, H., & Deb, K. (2015). U-nsga-iii: A unified evolutionary optimization procedure for single, multiple, and many objectives: Proof-of-principle results (A. Gaspar-Cunha, C. Henggeler Antunes, & C. C. Coello, Eds.). In A. Gaspar-Cunha, C. Henggeler Antunes, & C. C. Coello (Eds.), *Evolutionary multi-criterion optimization*, Cham, Springer International Publishing.

Stanton, A., & Channon, A. (2013). Heterogeneous complexification strategies robustly outperform homogeneous strategies for incremental evolution, In *Proceedings of the 12th european conference on artificial life*, Taormina, Italy, MIT Press.

Vargas, P. A., Paolo, E. A. D., Harvey, I., & Husbands, P. (Eds.). (2014). *The horizons of evolutionary robotics*. MIT Press.

Wagner, A. P., Zaman, L., Dworkin, I., & Ofria, C. (2020). Behavioral strategy chases promote the evolution of prey intelligence*. In W. Banzhaf, B. H. Cheng, K. Deb, K. E. Holekamp, R. E. Lenski, C. Ofria, R. T. Pennock, W. F. Punch, & D. J. Whittaker (Eds.), *Evolution in action: Past, present and future: A festschrift in honor of erik d. goodman* (pp. 225–246). Springer International Publishing. https://doi.org/10.1007/978-3-030-39831-6_17