



Productive fitness in diversity-aware evolutionary algorithms

Thomas Gabor¹ · Thomy Phan¹ · Claudia Linnhoff-Popien¹

Accepted: 15 March 2021 / Published online: 29 April 2021
© The Author(s) 2021

Abstract

In evolutionary algorithms, the notion of diversity has been adopted from biology and is used to describe the distribution of a population of solution candidates. While it has been known that maintaining a reasonable amount of diversity often benefits the overall result of the evolutionary optimization process by adjusting the exploration/exploitation trade-off, little has been known about what diversity is optimal. We introduce the notion of productive fitness based on the effect that a specific solution candidate has some generations down the evolutionary path. We derive the notion of final productive fitness, which is the ideal target fitness for any evolutionary process. Although it is inefficient to compute, we show empirically that it allows for an *a posteriori* analysis of how well a given evolutionary optimization process hit the ideal exploration/exploitation trade-off, providing insight into *why* diversity-aware evolutionary optimization often performs better.

Keywords Evolutionary algorithm · Adaptive fitness · Diversity

1 Introduction

Evolutionary algorithms are a widely used type of stochastic optimization that mimics biological evolution in nature. Like any other metaheuristic optimization algorithm (Brown et al. 2005; Conti et al. 2018), they need to maintain a balance on the exploration/exploitation trade-off in their search process: High exploration bears the risk to miss out on optimizing the intermediate solutions to the fullest; high exploitation bears the risk to miss the global optimum and get stuck in a sub-optimal part of the search space. Analogous to biological evolution, diversity within the population of solution candidates has been identified as a central feature to adjust the exploration/exploitation trade-off. Many means to maintain the diversity of the population throughout the process of evolution have been

developed in literature; comprehensive overviews are provided by Squillero and Tonda (2016) and Gabor et al. (2018), for example.

For problems with complex fitness landscapes, it is well known that increased exploration (via increased diversity) yields better overall results in the optimization, even when disregarding any diversity goal in the final evaluation (Ursem 2002; Toffolo and Benini 2003). However, this gives rise to a curious phenomenon: By augmenting the fitness function and thus making it match the original objective function *less*, we actually get results that optimize the original objective function *more*. This implies that any evolutionary algorithm does not immediately optimize for the fitness function it uses (but instead optimizes for a slightly different implicit goal). Furthermore, to really optimize for a given objective function, one should ideally use a (slightly) different fitness function for evolution. In this paper, we introduce *final productive fitness* as a theoretical approach to derive the ideal fitness function from a given objective function.

We see that final productive fitness cannot feasibly be computed in advance. However, we show how to approximate it *a posteriori*, i.e., when the optimization process is already finished. We show that the notion of final productive fitness is sound by applying it to the special case of diversity-aware evolutionary algorithms, which (for our

✉ Thomas Gabor
thomas.gabor@ifi.lmu.de

Thomy Phan
thomy.phan@ifi.lmu.de

Claudia Linnhoff-Popien
linnhoff@ifi.lmu.de

¹ LMU Munich, Oettingenstraße 67, 80538 München, Germany

purposes) are algorithms that directly encode a strife for increased diversity by altering the fitness of the individuals. By running these on various benchmark problems, we empirically show that diversity-aware evolutionary processes might just approximate final productive fitness more accurately than an evolutionary process using just the original objective. We show that the fitness alteration performed by these algorithms, when it improves overall performance, does so *while* (perhaps *because*) it better approximates final productive fitness. We thus argue that the notion of final productive fitness for the first time provides a model of *how* diversity is beneficial to evolutionary optimization, which has been called for by various works in literature:

- “One of the urgent steps for future research work is to better understand the influence of diversity for achieving good balance between exploration and exploitation.” (Črepinšek et al. 2013),
- “This tendency to discover both quality and diversity at the same time differs from many of the conventional algorithms of machine learning, and also thereby suggests a different foundation for inferring the approach of greatest potential for evolutionary algorithms.” (Pugh et al. 2016),
- “However, the fragmentation of the field and the difference in terminology led to a general dispersion of this important corpus of knowledge in many small, hard-to-track research lines” and, “[w]hile diversity preservation is essential, the main challenge for scholars is devising general methodologies that could be applied seamlessly [...]” (Squillero and Tonda 2016).

It should be noted that the approach presented in this paper merely provides a new perspective on exploration/exploitation in evolutionary algorithms and a new method of analyzing the effects of diversity. It is up to future works to derive new means to actively promote diversity from this analysis.

In this paper, we provide a short mathematical description of evolutionary processes in Sect. 2 and build our notion of (final) productive fitness on top of that in Sect. 3. Section 4 describes the empirical results and Sect. 5 discusses related work before Sect. 6 concludes.

2 Foundations

For this paper, we assume an evolutionary process (EP) to be defined as follows: Given a fitness function $f : \mathcal{X} \rightarrow [0; 1] \subset \mathbb{R}$ for an arbitrary set \mathcal{X} called the search space, we want to find an individual $x \in \mathcal{X}$ with the best fitness $f(x)$. For a maximization problem, the best fitness is that of an individual x so that $f(x) \geq f(x') \forall x' \in \mathcal{X}$. For a

minimization problem, the best fitness is that of an individual x so that $f(x) \leq f(x') \forall x' \in \mathcal{X}$. Note that we normalize our fitness space on $[0; 1] \subset \mathbb{R}$ for all problems for ease of comparison. Whenever the maximum and minimum fitness are bounded, this can be done without loss of generality.

Usually, the search space \mathcal{X} is too large or too complicated to guarantee that we can find the exact best individual(s) using standard computing models (and physically realistic time). Thus, we take discrete subsets of the search space \mathcal{X} via sampling and iteratively improve their fitness. An evolutionary process \mathcal{E} over g generations, $g \in \mathbb{N}$, is defined as $\mathcal{E} = \langle \mathcal{X}, e, f, (X_i)_{i < g} \rangle$. \mathcal{X} is the search space. $e : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$ is the evolutionary step function so that $X_{i+1} = e(X_i) \forall i \geq 0$. As defined above, $f : \mathcal{X} \rightarrow [0; 1] \subset \mathbb{R}$ is the fitness function. $(X_i)_{i < g}$ is a series of populations so that $X_i \subseteq \mathcal{X} \forall i$ and X_0 is the initial population. Note that as the evolutionary step function e is usually non-deterministic, we define $E(X) = \{X' | X' = e(X)\}$ to be the set of all possible next populations.

We use the following evolutionary operators:

- The recombination operator $rec : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$ generates a new individual from two individuals.
- The mutation operator $mut : \mathcal{X} \rightarrow \mathcal{X}$ alters a given individual slightly to return a new one.
- The migration operator $mig : \mathcal{X}$ generates a random individual $mig() \in \mathcal{X}$.
- The (survivors) selection operator $sel : 2^{\mathcal{X}} \times \mathbb{N} \rightarrow 2^{\mathcal{X}}$ returns a new population $X' = sel(X, n)$ given a population $X \subseteq \mathcal{X}$, so that $|X'| \leq n$.

The operators rec, mut, mig can be applied to a population X by choosing individuals from X to fill their parameters (if any) according to some selection scheme $\sigma : 2^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$ and adding their return to the population. For example, we allow to write $mut_{\sigma}(X) = X \cup \{mut(x') \mid x' \in \sigma(X)\}$. Note that all children are added to the population and do not replace their parents in this formulation.

For any evolutionary process $\mathcal{E} = \langle \mathcal{X}, e, f, (X_i)_{i < g} \rangle$ and selection schemes $\sigma_1, \sigma_2, \sigma_3$ we assume that

$$X_{i+1} = e(X_i) = sel(mig_{\sigma_3}(mut_{\sigma_2}(rec_{\sigma_1}(X_i))), |X_i|). \quad (1)$$

Usually, we assume that an evolutionary process fulfills its purpose if the best fitness of the population tends to become better over time, i.e., given a sufficiently large amount of generations $k \in \mathbb{N}$, it holds for maximization problems that $\max_{x \in X_i} f(x) < \max_{x \in X_{i+k}} f(x)$. We define the overall result of an evolutionary process $\mathcal{E} = \langle \mathcal{X}, e, f, (X_i)_{i < g} \rangle$ with respect to a fitness function ϕ (which may or may not be different from the fitness f used during evolution) to be best value found and kept in evolution, i.e., for a maximizing objective ϕ we define

$$|\mathcal{E}|_\phi = \max_{x \in X_g} \phi(x). \quad (2)$$

Note that there are evolutionary processes which include a *hall-of-fame* mechanism, i.e., are able to return the result fitness

$$||\mathcal{E}||_\phi = \max_{i=1, \dots, g} \max_{x \in X_i} \phi(x). \quad (3)$$

However, we can derive the equality $|\mathcal{E}|_\phi = ||\mathcal{E}||_\phi$ when we assume *elitism* with respect to ϕ , i.e., $\arg \max_{x \in X_i} \phi(x) \in X_{i+1}$ for all $i = 1, \dots, g$. Since it makes reasoning easier and hardly comes with any drawback for sufficiently large populations, we use elitist evolutionary processes (with respect to f) from here on.

3 Approach

The central observation we build our analysis on is that in many cases the results of optimizing for a given objective function (called OF) can be improved by not using OF as a the fitness function f of the evolutionary process directly. Consequently, changing the fitness function f away from the true objective OF in some cases leads to better results with respect to the original objective function OF . Note that this phenomenon extends beyond just heuristic optimization and is known as *reward shaping* in reinforcement learning, for example (Ng et al. 1999).

In evolutionary algorithms oftentimes a property called diversity is considered in addition to the objective function OF to improve the progress of the evolutionary process (Gabor et al. 2018; Squillero and Tonda 2016; Ursem 2002). In some way or the other, diversity-enhancing evolutionary algorithms award individuals of the population for being different from other individuals in the population. While there are many ways to implement this behavior, like topology-based methods (Tomassini 2006), fitness sharing (Sareni and Krahenbuhl 1998), ensembling (Hart and Sim 2018), etc., we consider an instance of diversity-enhancing evolutionary algorithms that is simpler to analyze: By quantifying the distance of a single individual to the population, we can define a secondary fitness SF that rewards high diversity in the individual. This approach was shown by Wineberg and Oppacher (2003) to be an adequate general representation of most well-known means of measuring diversity in a population.

In order to avoid the difficulties of multi-objective evolution, we can then define the augmented fitness function AF that incorporates both the objective fitness OF and the secondary fitness SF into one fitness function to be used for the evolutionary process.

Definition 1 (Augmented Fitness) Given the objective fitness OF , a diversity-aware secondary fitness SF , and a diversity weight $\lambda \in [0; 1] \subset \mathbb{R}$, we define the augmented fitness AF as

$$\text{AF}(x) = (1 - \lambda) \cdot \text{OF}(x) + \lambda \cdot \text{SF}(x). \quad (4)$$

As is shown in Gabor et al. (2018) and Wineberg and Oppacher (2003) such a definition of the augmented fitness suffices to show benefits of employing diversity.

We can then define two evolutionary processes $\mathcal{E}_{\text{OF}} = \langle \mathcal{X}, e, \text{OF}, (X_i)_{i < g} \rangle$ and $\mathcal{E}_{\text{AF}} = \langle \mathcal{X}, e, \text{AF}, (X'_i)_{i < g} \rangle$. We observe the curious phenomenon that in many cases the augmented fitness AF better optimizes for OF than using OF itself, formally

$$|\mathcal{E}_{\text{OF}}|_{\text{OF}} < |\mathcal{E}_{\text{AF}}|_{\text{OF}}, \quad (5)$$

which raises the following question: If OF is not the ideal fitness function to optimize for the objective OF , what is?

Given a sequence of populations $(X_i)_{i < g}$ spanning over multiple generations $i = 1, \dots, g$ we can write down what we actually want our population to be like inductively starting from the last generation g : The net benefit of X_g to our (maximizing) optimization process is exactly

$$|\mathcal{E}|_{\text{OF}} = \max_{x \in X_g} \text{OF}(x) \quad (6)$$

as this population will not evolve any further and thus the best individual within X_g is what we are going to be stuck with as the result of the optimization process.

Note that the individuals of X_{g-1} already contribute differently to the result of the optimization process: From the perspective of generation $g - 1$ the overall optimization result is

$$\max_{x \in X_{g-1}} \max_{x' \in X_g(x)} \text{OF}(x') \quad (7)$$

where the follow-up generation $X_g(x)$ is any¹ population from $\{X_g \mid X_g \in E(X_{g-1}) \wedge x \in X_g\}$, i.e., the possible next populations where x survived.

Intuitively, the contribution of the the second-to-last generation X_{g-1} to the result of the optimization process stems from the objective fitness OF that this generation's individuals can still achieve in the final generation X_g . Generally, this does not fully coincide with the application of the objective function OF in said generation:

$$\max_{x \in X_{g-1}} \max_{x' \in X_g(x)} \text{OF}(x') \neq \max_{x \in X_{g-1}} \text{OF}(x) \quad (8)$$

¹ Arguments can be made to pick either the average or the maximum over all these populations. We discuss both cases later in the text.

That means: While rating individuals according to their objective fitness OF in the last generation of the evolutionary process is adequate, the actual benefit of the individual x to the optimization result and the value of $OF(x)$ may diverge more the earlier we are in the evolutionary process. Accordingly, at the beginning of an evolutionary process, the objective fitness OF might not be a good estimate of how much the individuals will contribute to the process's return with respect to OF at the end of the optimization process. Still, standard optimization techniques often use the objective fitness OF as a (sole) guideline for the optimization process.

Instead, we ideally want to make every decision (mutation, recombination, survival, ...) at every generation X_i with the ideal result for the following generations X_{i+1}, X_{i+2}, \dots and ultimately the final generation X_g in mind. We call this the *optimal evolutionary process*. Obviously, to make the optimal decision early on, we would need to simulate all the way to the end of the evolution, including all the follow-up decisions. This renders optimal evolution infeasible as an algorithm. However, we can use it for a posteriori analysis of what has happened within a different evolutionary process. In order to do so, we need to give a fitness function for the optimal process (as it obviously should not be OF).

Instead, we formalize the benefit to the optimization process discussed above and thus introduce the notion of productive fitness. But first, we need a simple definition on the inter-generational relationships between individuals.

Definition 2 (Descendants) Given an individual x in the population of generation i , $x \in X_i$, of an evolutionary process. All individuals $x' \in X_{i+1}$ so that x' resulted from x via a mutation operator, i.e., $x' = \text{mut}(x)$, or a recombination operator with any other parent, i.e., there exists $y \in X_i$ so that $x' = \text{rec}(x, y)$, are called *direct descendants* of x . Further given a series of populations $(X_i)_{0 \leq i < g}$ we define the set of all descendants D_x as the transitive hull on all direct descendants of x .

We can now use this relationship to assign the benefit that a single individual has had to the evolution a posteriori. For this, we simply average the fitness of all its surviving descendants.

Definition 3 (Productive Fitness) Given an individual x in the population of generation i , $x \in X_i$, of an evolutionary process. Let $D_x \subseteq \mathcal{X}$ be the set of all descendants from x . The *productive fitness after n generations* or *n -productive fitness* is the average objective fitness of x 's descendants, written

$$PF_n(x) = \begin{cases} \text{avg}_{x' \in D_x \cap X_{i+n}} OF(x') & \text{if } D_x \cap X_{i+n} \neq \emptyset \\ w & \text{otherwise.} \end{cases} \quad (9)$$

Note that in case the individual x has no descendants in n generations, we set its productive fitness $PF_n(x)$ to a worst case value w , which in our case of bounded fitness values is 0 for maximizing optimization processes and 1 for minimizing optimization processes.

We argue that the productive fitness PF is better able to describe the actual benefit the individual brings to the optimization process, as represented by what parts of the individual still remain inside the population in a few generations. Note that our notion of productive fitness is rather harsh in two points:

- We only take the average of all descendants' fitness. One could argue that we may want a more optimistic approach where we might reward the individual for the best offspring it has given rise to. However, we argue that every bad individual binds additional resources for eliminating it down the road and thus a low target accuracy should actively be discouraged.
- When the line of an individual dies out completely, we assign the worst possible fitness. Arguments could be made that even dead lines contribute to the search process by ruling out unpromising areas while, e.g., increasing the diversity scores of individuals in more promising areas of the search space. Still, we do count any however distant descendants, so even small contributions to the final population avoid the penalty w .

We leave the analysis of the effects of the discussed parameters to future work. Note that for now, our notion of productive fitness only covers a fixed horizon into the future. We can trivially extend this definition to respect the final generation no matter what generation the current individual is from:

Definition 4 (Final Productive Fitness) Given an individual x in the population of generation i , $x \in X_i$, of an evolutionary process of g generations in total. The final productive fitness of x is the fitness of its descendants in the final generation, i.e., $FPF(x) = PF_{g-i}(x)$.

We argue that final productive fitness is able to describe what the fitness function of an optimal evolutionary process looks like: Every evaluation is done in regard to the contribution to the final generation, i.e., the ultimate solution returned by the search process.

Thesis 1 When rolling the ideal choices in all randomized evolutionary operators, final productive fitness FPF is the optimal fitness function for evolutionary processes, i.e., an evolutionary process yields the best results when it optimizes for FPF at every generation.

We sketch a short argument in favor of Thesis 1. For a more in-depth discussion, see Gabor and Linnhoff-Popien

(2020). Let $\mathcal{E}_{\text{FPF}} = \langle \mathcal{X}, e, \text{FPF}, (X_i^{\text{FPF}})_{i < g} \rangle$ be an evolutionary process using final productive fitness FPF . Let $\mathcal{E}_{\text{IDF}} = \langle \mathcal{X}, e, \text{IDF}, (X_i^{\text{IDF}})_{i < g} \rangle$ be an evolutionary process using a different (possibly more ideal) fitness IDF . Let $X_0^{\text{FPF}} = X_0^{\text{IDF}}$. We assume that

$$\max_{x \in X_g^{\text{FPF}}} \text{OF}(x) < \max_{x \in X_g^{\text{IDF}}} \text{OF}(x). \quad (10)$$

Since Eq. 10 implies that at least $X_g^{\text{FPF}} \neq X_g^{\text{IDF}}$, there is an individual $x \in X_g^{\text{IDF}}$ so that $x \notin X_g^{\text{FPF}}$ and $\text{OF}(x) > \max_{y \in X_g^{\text{FPF}}} \text{OF}(y)$. Since both \mathcal{E}_{FPF} and \mathcal{E}_{IDF} use the same evolutionary step function e except for the used fitness, their difference regarding x needs to stem from the fact that there exists an individual x' that is an ancestor of x , i.e., $x \in D_{x'}$, so that x' was selected for survival in \mathcal{E}_{IDF} and not in \mathcal{E}_{FPF} , which implies that $\text{FPF}(x') < \text{IDF}(x')$. However, since x is a possible descendant for x' , the computation of $\text{FPF}(x')$ should have taken $\text{OF}(x)$ into account,² meaning that x' should have survived in \mathcal{E}_{FPF} after all, which contradicts the previous assumption. \square

Of course, Thesis 1 is a purely theoretical argument as we cannot guarantee optimal choices in usually randomized evolutionary operators and productive fitness in general thus comes with the reasonable disadvantage that it cannot be fully computed in advance. But for a given, completed run of an evolutionary process, we can compute the factual FPF single individuals had a posteriori. There, we still do not make optimal random choices but just assume the ones made as given.

Still, we take Thesis 1 as hint that final productive fitness might be the right target to strive for. We argue that augmenting the objective fitness OF (even with easily computable secondary fitness functions) may result in a fitness function which better approximates final productive fitness FPF . In the following Sect. 4, we show empirically that (in the instances where it helps³) diversity-based secondary fitness SF resembles the final productive fitness FPF of individuals much better than the raw objective function OF does.

Thesis 2 When a diversity-aware augmented fitness function AF is aiding the evolutionary optimization process with respect to an objective fitness OF , it is doing so by approximating the final productive fitness FPF of a converged evolutionary process in a more stable way (i.e.,

more closely when disregarding the respective scaling of the fitness functions) throughout the generations of the evolutionary process.

This connection not only explains why diversity-aware fitness functions fare better than the pure objective fitness but also poses a first step towards a description how to deliberately construct diversity-aware fitness functions, knowing that their ideal purpose is to approximate the not fully computable final productive fitness. Again, we refer to Gabor and Linnhoff-Popien (2020) for more elaborate theoretical arguments.

Since we cannot estimate all possible futures for an evolutionary process, we provide empirical evidence in favor of Thesis 2 using a a posteriori approximation: Given an already finished evolutionary process, we compute the FPF values given only those individuals that actually came into being during that single evolutionary process (instead of using all possible descendants). We argue that this approximation is valid because if the evolutionary process was somewhat successful, then all individuals' descendants should be somewhat close to their ideal descendants *most likely*.⁴ Note that the reverse property is not true (i.e., even in a bad run, individuals still aim to generate better descendants, not worse), which is why our approximation does not permit any statements about augmented fitness that does not aid the evolutionary process.

4 Experiments

For all experiments, we run an evolutionary process as defined in Sect. 2 with a mutation operator *mut* that adds a (possibly negative) random value to one dimension of individual, applied with rate 0.1 to all individuals at random. For *rec* we apply random crossover with rate 0.3 for a single individual and a randomly chosen mate. We apply *mig* with a rate of 0.1 (Gabor et al. 2018). Following Wineberg and Oppacher (2003) and the results in Gabor et al. (2018), we focus on a Manhattan distance function for the secondary fitness; we also plot evolutionary processes using fitness sharing with parameter $\alpha = 2.0$ and dissimilarity threshold $\sigma = n$, where n is the dimensionality of the problem (Sareni and Krahenbuhl 1998), or inherited fitness with inheritance weight $\kappa = 0.5$ (Chen et al. 2002; Gabor et al. 2018) for comparison,

² Note that $\text{OF}(x)$ cannot be compensated by other descendants of x' with possibly bad objective fitness even as we average the results because all offspring is created by a random choice, which we assume to be ideal. This also shows how strong that assumption is.

³ Note the gravity of that restriction: We do not consider failed runs of evolutionary algorithms since we have no assumptions on how FPF should behave, i.e., relate to AF , in that case. Future work may fill that void.

⁴ Note that we could construct a terrible evolutionary process that just happens to find the global optimum in the last generation out of the blue via random migration. That process would have a poor stability between AF and FPF but a very successful result. However, since evolution at every step tries not to be terrible, we consider that scenario to be quite unlikely so that it should not play a role when we analyze the augmented fitness on multiple runs, parametrizations, and domains.

since both approaches also use an adapted fitness function to promote diversity.⁵ The selection operator *sel* is a simple rank-based cut-off in the shown evolutionary processes. Cut-off with protection for new individuals as well as roulette wheel selection was also tested without yielding noticeably different results.

All code that produced the results of this paper is available at github.com/thomasgabor/naco-evolib.

4.1 Pathfinding

We start with the pathfinding problem, which was shown to greatly benefit from employing diversity in the optimization process (Gabor et al. 2018): Given a room of dimensions 1×1 , we imagine a robot standing at position $(0.5, 0.1)$. It needs to reach a target area at the opposite side of the room. See Fig. 1 for an illustration. The room also features a huge obstacle in the middle and thus the robot needs to decide on a way around it. The agent can move by performing an action $a \in \{(\delta x, \delta y) \mid -0.33 < \delta x < 0.33, -0.33 < \delta y < 0.33\}$. A single solution candidate consists of $n = 5$ actions $\langle (\delta x_i, \delta y_i) \rangle_{1 \leq i \leq n}$. It achieves a reward of $\frac{1}{n} = 0.2$ every time it stays within the target area between steps, i.e., its fitness is given via $OF(\langle (\delta x_i, \delta y_i) \rangle_{1 \leq i \leq n}) = r(\langle (\delta x_i, \delta y_i) \rangle_{1 \leq i \leq n}, (0.5, 0.1))$ where

$$r(\langle (\delta x_i, \delta y_i) \rangle_{1 \leq i \leq n}, (x, y)) = r(\langle (\delta x_i, \delta y_i) \rangle_{2 \leq i \leq n}, (x + \delta x_1, y + \delta y_1)) + t(x, y) \quad (11)$$

$$\text{with } r(\langle \rangle, (x, y)) = t(x, y) \quad (12)$$

$$\text{and } t(x, y) = \begin{cases} \frac{1}{n} & \text{if } 0.4 \leq x \leq 0.6 \wedge 0.8 \leq y \leq 1.0 \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The pathfinding problem lends itself to the application of diversity, as the optimization process in most cases first strikes a local optimum where it reaches the target area sometime by accident (and most probably towards the end of its steps). It then needs to switch to the global optimum where the first three steps are as goal-directed as possible and the last two steps are very small in order to stay within the target area.

We now compare a standard evolutionary algorithm given only the objective function $OF(x) = f(x)$ to a diversity-aware evolutionary algorithm using the Manhattan distance on the solution candidate structure as a secondary




Fig. 1 The Pathfinding problem. A robot (red) needs to find a fixed path to reach the target area (green)

fitness function, i.e., $AF(x) = (1 - \lambda) \cdot OF(x) + \lambda \cdot SF(x)$ as given in Definition 1

$$\text{where } SF(x) = \frac{1}{2n} \cdot \text{avg}_{x' \in \sigma_4(X)} \text{manhattan}(x, x') \quad (14)$$

and $\text{manhattan}(\langle (\delta x_i, \delta y_i) \rangle_{1 \leq i \leq n},$

$$\langle (\delta x'_i, \delta y'_i) \rangle_{1 \leq i \leq m}) = \sum_{i=1}^m |\delta x_i - \delta x'_i| + |\delta y_i - \delta y'_i|. \quad (15)$$

Note that σ_4 is a selection function that randomly selects 10 individuals from the population X . We use it to reduce the computational cost of computing the pairwise distance for all individuals in the population. Its admissibility for approximating the full pairwise distance was shown in Gabor and Belzner (2017). Just as we normalized the fitness function f to $[0; 1] \subset \mathbb{R}$ we also normalize the secondary fitness sf to the same range via division by the maximum Manhattan distance between two individuals, i.e., $2n$, to make the combination easier to understand. For now, we set $\lambda = 0.4$, which we discuss later.

Each evolution was run 30 times for 1500 generations each, using a population size of 50. Figure 2a shows the best fitness achieved per generation for all tested approaches. We see that (especially distance-based) diversity-aware evolution produces much better objective results. Figure 2b shows the separate diversity score sf maintained by the best individual, which can only be computed in a meaningful way for Manhattan diversity. In Fig. 2c the standard approach shows the same plot as before since its fitness is not augmented. For all other approaches we plot

⁵ Of course, many more diversity-aware evolutionary algorithms could have been analyzed here. We would like to refer to our experiments in Gabor et al. (2018) as well as the survey by Squillero and Tonda (2016) for a comprehensive overview.

the augmented fitness AF that is actually used for selection. We see that due to the combination of distance and objective fitness, Manhattan-diverse evolution starts higher but climbs slower than the respective objective fitness. Fitness sharing results in very small absolute values for fitness but climbs up nonetheless.

From the already run evolutionary processes, we can compute the final productive fitness as given in Definitions 3 and 4 a posteriori. Figure 2d shows the maximum FPF per generation. We see that Manhattan-diverse and inherited fitness maintain a rather continuous lineage from the initial population to the best solution in the final generation as the final fitness propagates to the final productive fitness of very early generations. This behavior is rather unsurprising but illustrates the notion of the final productive fitness that measures the individuals' impact in the final result.

For Fig. 2e we compute the perhaps most interesting measurement: This plot shows for each population X in a given generation the result of $\text{avg}_{x \in X} |AF(x) - FPF(x)|$, i.e., the average difference between the augmented fitness and the final productive fitness *per individual*. Thus, we get to assess how well the augmented fitness approximates the final productive fitness. There are a few observations to be made:

1. Towards the last few generations, we notice a rapid spike in the FPF as the amount of descendants in the final generation to be considered for the FPF decreases fast.
2. The actual value of the distance (i.e., the height of the line) is irrelevant to the analysis of Thesis 2 and largely determined by the setting of λ .
3. Throughout most of the plot, the Manhattan-diverse evolution maintains a relatively stable level, i.e., the augmented fitness AF approximates the final productive fitness FPF throughout the evolution. The less stable evolutions also show a worse overall result.

To further elaborate on that last point, we consider Fig. 2: It shows the average absolute value of change over 150-generations-wide windows of the $|AF(x) - FPF(x)|$ metric used in Fig. 2e. The plot was smoothed using a convolution kernel $\langle 1, \dots, 1 \rangle$ of size 25. Roughly speaking, we can see the slope of the plots in Fig. 2e here. In this plot, good evolutionary processes should maintain rather low values according to Thesis 2. We can observe that Manhattan-diverse evolution maintains the lowest values almost throughout the entire evolution. While fitness sharing shows increases and decreases in matching the FPF at a higher level than Manhattan diversity, inherited fitness shows a huge spike in the beginning (as does the standard approach), thus making a much less stable match for the FPF. As proposed by Thesis 2, the match between AF and FPF

roughly corresponds to the quality of the overall result of the evolutionary process.

As mentioned earlier, we also further analyzed the importance of the setting for λ for the evolution. Figure 3 shows the impact of λ on the best results generated by the evolution. $\lambda = 0$ equals the standard evolution in all previous plots. Unsurprisingly, we see that some amount of diversity-awareness improves the results of evolution but setting λ too high favors diversity over the actual objective fitness and thus yields very bad results. We want to add that more intricate version of Manhattan-based augmented fitness AF might aim to adjust the λ parameter during evolution just as inherited fitness and fitness sharing might want to adjust their parameters. For these experiments, we chose a static parameter setting for simplicity.

4.2 The route planning problem

The route planning problem is a discrete optimization problem with a similar motivation as the pathfinding problem. Again, we adapt the problem and its description from Gabor et al. (2018).

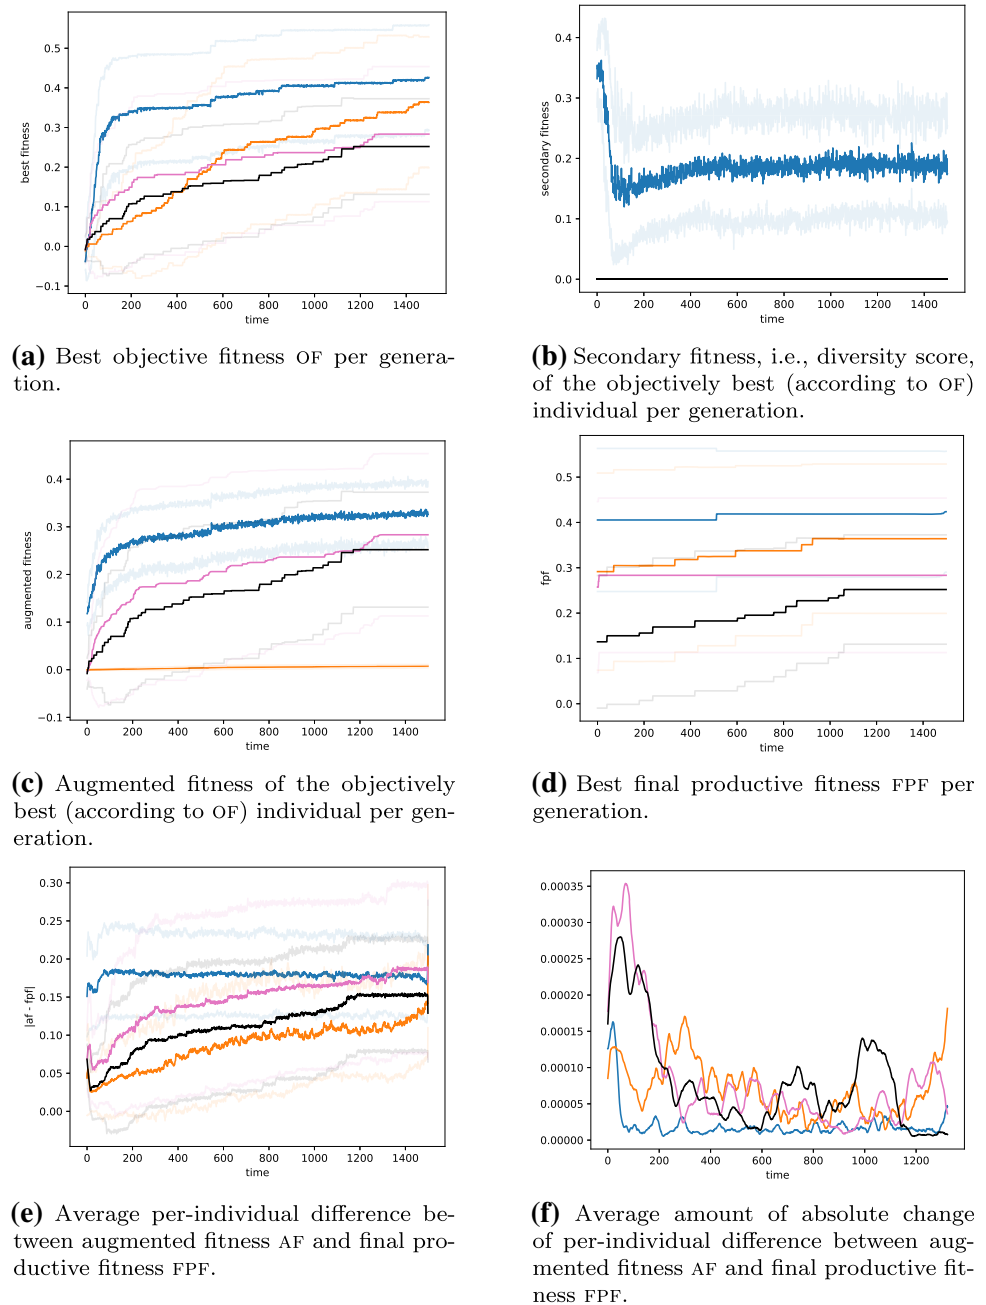
A robot needs to perform $n = 12$ different tasks in a fixed order by visiting relevant workstations. Each workstation can perform exactly one of the tasks and for each task, there are $o = 5$ identical workstations to choose from. Accordingly, a solution candidate is a vector $\langle w_1, \dots, w_n \rangle$ with $w_i \in \{1, \dots, o\}$ for all $1 \leq i \leq n$. See Fig. 4 for an illustration using a smaller setting. A single workstation W can be identified by a tuple of its task type and its number, i.e., $W = (i, k)$ for some $1 \leq i \leq n$ and $1 \leq k \leq o$. To mimic various means of transport, the distance $D(W, W')$ between every two workstations $W = (i, k)$ and $W' = (j, l)$ is randomized individually within a range of $[0, \frac{1}{n}] \subset \mathbb{R}$. Note that this (most likely) gives rise to a non-euclidean space the robot is navigating. The objective fitness for this minimization problem is given via

$$\begin{aligned} \text{OF}(\langle w_1, \dots, w_n \rangle) \\ = \sum_{i=1}^{n-1} D((i, w_i), (i+1, w_{i+1})). \end{aligned} \quad (16)$$

Again, we from here also construct an augmented fitness $AF(x) = (1 - \lambda) \cdot \text{OF}(x) + \lambda \cdot \text{SF}(x)$ (cf. Definition 1) but now use the Hamming distance as a secondary fitness so that

$$\text{SF}(x) = \frac{1}{2n} \cdot \text{avg}_{x' \in \sigma_4(X)} \text{hamming}(x, x') \quad (17)$$

Fig. 2 Evolution for the pathfinding problem. Standard evolutionary process using OF shown in black, diversity-aware evolutionary process using AF with Manhattan distance shown in blue. Inherited fitness (purple) and fitness sharing (orange) shown for comparison. All results averaged over 25 independent runs, the standard deviation is shown in transparent lines



$$\text{where } \text{hamming}(\langle w_1, \dots, w_n \rangle, \langle w'_1, \dots, w'_n \rangle) = \sum_{i=1}^n h(w_i, w'_i) \quad (18)$$

$$\text{and } h(w, w') = \begin{cases} 0 & \text{if } w = w' \\ 1 & \text{otherwise.} \end{cases} \quad (19)$$

Besides, we apply the same evolutionary processes as in Sect. 4.1 but the parameter search shown in Fig. 5 now recommended $\lambda = 0.25$ for weighting now Hamming-based diversity.⁶ We evolve 20 independent populations of size 50 for 400 generations and plot the same data we have

seen before: Fig. 6a shows the best fitness achieved in evolution. Inherited fitness takes a lot more time but eventually almost reaches the level of Manhattan-diversity. However, both methods yield similarly solid results as fitness sharing or the naïve algorithm. This is mirrored by all methods showing quite stable behavior in Figs. 6e and 6f with the standard approach showing the highest fall within the first few generations as it matches FPF the least.

⁶ It should be noted that the difference for various settings of diversity weights (including $\lambda = 0$) is much less pronounced for this domain as can be seen in Fig. 5.




Fig. 3 Parameter analysis for the diversity weight λ . We show best fitness among all generations for 20 different settings of $\lambda = 0.0, 0.05, \dots, 1.0$. All results averaged over 25 independent runs each, the standard deviation shown in transparent lines




Fig. 4 Illustration of the route planning problem for $n = 3$ tasks and $o = 5$ workstations per task

However, the results for all evolutions are very close together for the route planning problem.

4.3 Schwefel

Finally, we consider one⁷ of the canonical benchmark problems for evolutionary algorithms. The implementation of the Schwefel problem is taken from Rainville et al. (2012) while our study on the impact of diversity again follows experiments performed in Gabor et al. (2018).

The original fitness function is given as

⁷ We also tested other famous problems like Rainville et al. (2012) with essentially similar results.




Fig. 5 Parameter analysis for the diversity weight λ for the route planning problem. We show best fitness among all generations for 20 different settings of $\lambda = 0.0, 0.05, \dots, 1.0$. All results averaged over 20 independent runs each, the standard deviation shown in transparent lines

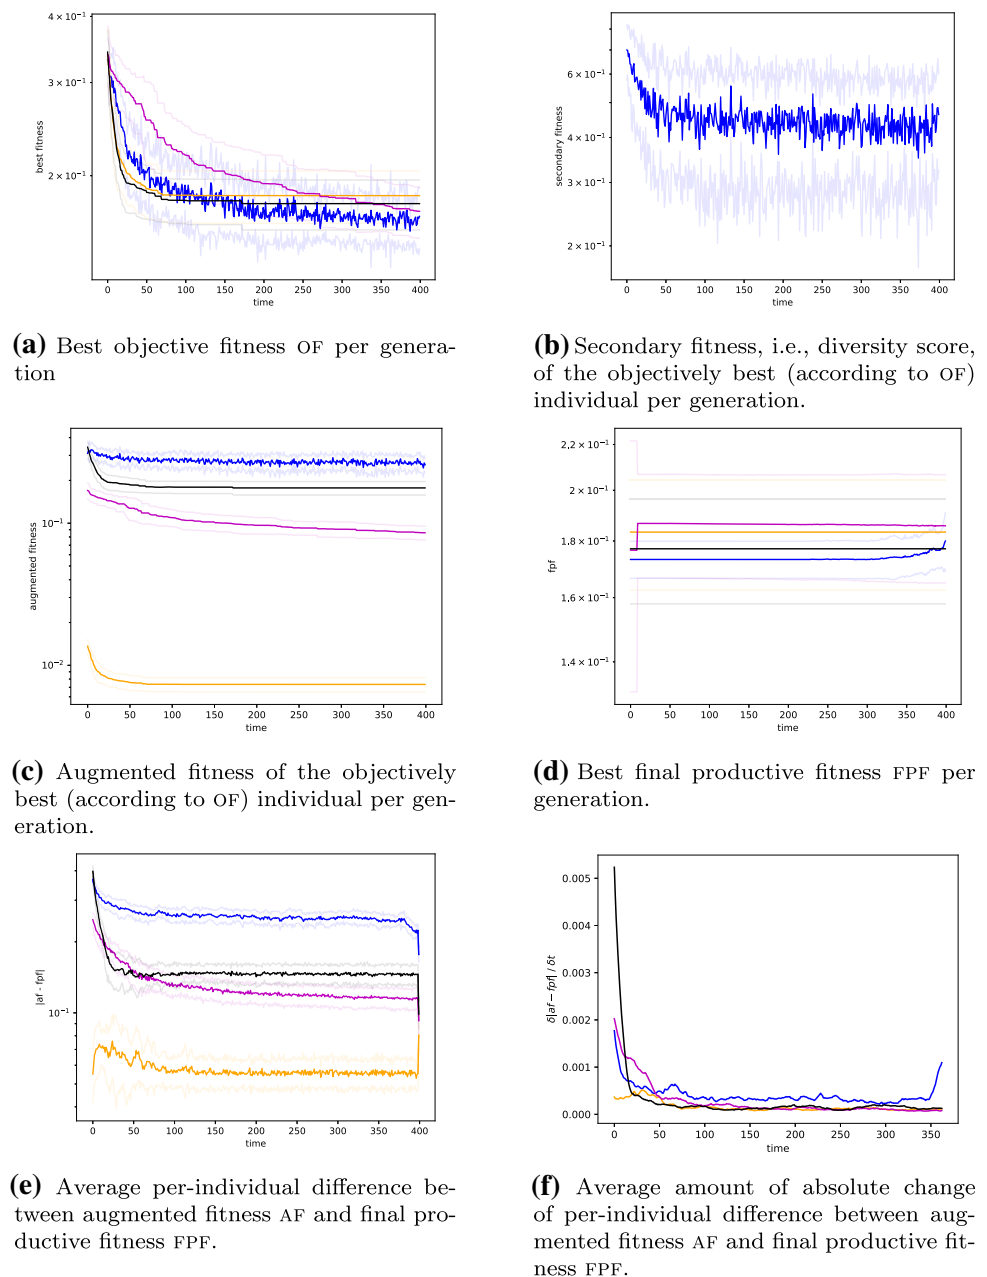
$$\text{schwefel}(\langle x_1, \dots, x_n \rangle) = 418.9828872724339 \cdot n - \sum_{i=1}^n x_i \cdot \sin(\sqrt{|x_i|}) \quad (20)$$

with $x_1, \dots, x_n \in [-500; 500] \subset \mathbb{R}$ where $n = 8$ is the dimensionality we use in our experiments (Rainville et al. 2012).⁸ The resulting function is illustrated in Fig. 7. The Schwefel problem is a minimization problem looking for an x so that $\text{schwefel}(x) = 0$. Again, we normalize the result values defining $\text{OF}(x) = \frac{1}{4000} \cdot \text{schwefel}(x)$. Manhattan distance uses diversity weight $\lambda = 0.3$ as suggested by Fig. 8.

We run the same kind of analysis as for the previous problems and plot the same data in Fig. 9. These experiments were performed on populations of size 50 evolving for 400 generations. Figure 9a shows a mixed picture: The Manhattan-diverse evolution again outperforms the standard approach, but inherited fitness hardly yields any benefit while fitness sharing performs worse than the standard approach. In Fig. 9d we see a different picture than before: All final productive fitness values are not as stable anymore but vary throughout the evolution, suggesting that solutions to the Schwefel problem are more influenced by migration (and thus show no continuous lineage) than for the other problems considered. Fig. 9e shows the $|\text{AF} - \text{FPF}|$ metric, which measures the individual difference between augmented fitness and final productive fitness. We observe very clearly that, after a short starting

⁸ As can be seen in Eq. 20, the Schwefel function is identical in each dimension. Choosing a higher-dimensional instance thus does not substantially change the nature of the fitness landscape. For ease of analysis, we settle for $n = 8$.

Fig. 6 Evolution for the route planning problem. Standard evolutionary process using OF shown in black, diversity-aware evolutionary process using AF with Hamming distance shown in blue. Inherited fitness (purple) and fitness sharing (orange) shown for comparison. Results averaged over 20 independent runs, the standard deviation is shown in transparent lines



phase, this distance remains much more stable for the Manhattan-diverse evolution than any other approach, indicating that the augmented fitness is approximating the final productive fitness. Note again that in the last few generations, computing the final productive fitness has limited meaning. In Fig. 9f this behavior shows clearly as Manhattan-diverse evolution forms a line at the very bottom of the plot with almost no change in how AF matches FPF. All other approaches, which perform noticeably worse, also show a much more erratic pattern in how well their augmented fitness matches the final productive fitness.

5 Related work

Diversity has been a central topic of research in evolutionary algorithms (den Heijer and Eiben 2012; Morrison and De Jong 2001; Toffolo and Benini 2003; Ursem 2002). Its positive effect on the evolutionary process has often been observed there, but rarely been interpreted beyond a biological metaphor, i.e., “diversity is a key element of the biological theory of natural selection and maintaining high diversity is supposed to be generally beneficial” (Corno et al. 2005).




Fig. 7 Illustration of the Schwefel function for $n = 2$ dimensions. Image taken from (Benchmarks 2020)




Fig. 8 Parameter analysis for the diversity weight λ for the Schwefel problem. We show best fitness among all generations for 10 different settings of $\lambda = 0.0, 0.1, \dots, 1.0$. All results averaged over 20 independent runs each, the standard deviation shown in transparent lines

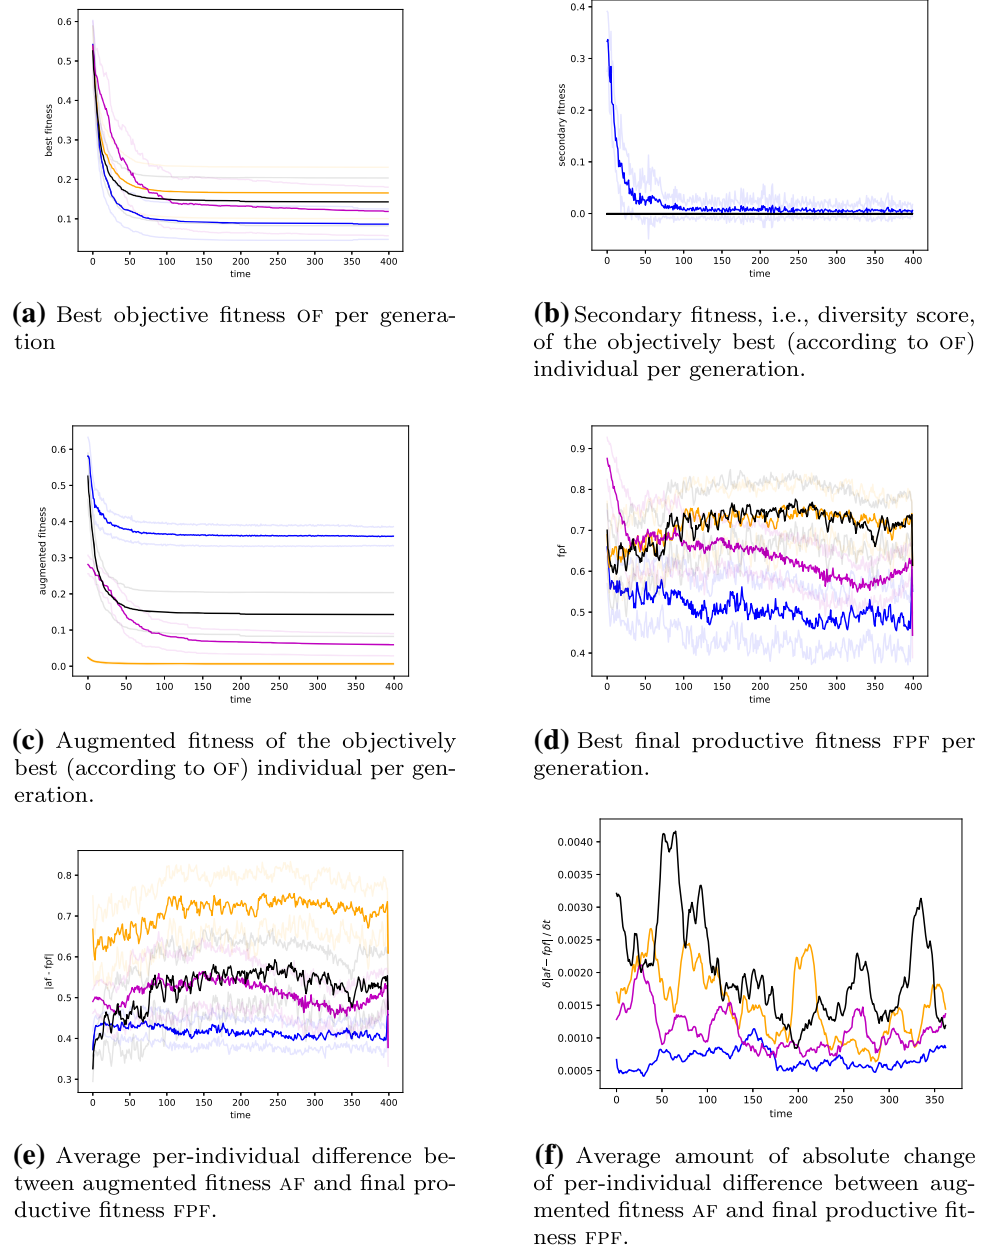
Without much concept of what to look for in a mechanism for diversity-awareness, lots of variants have spawned in research. Instead of repeating them, we would like to point out a few resources for a comprehensive overview: Burke et al. (2004), among others like Brameier and Banzhaf (2002 or McPhee and Hopper (1999) discuss various means to measure and promote diversity in genetic programming, which for the most part should apply to all evolutionary algorithms. They also provide an extensive analysis of the connection between diversity and achieved fitness, but do not define productive fitness or a similar notion. A more recent comprehensive overview of means

to describe and enable diversity has been put together by Squillero and Tonda (2016), also providing a taxonomy on various classes of approaches to diversity. Gabor et al. (2018) provide a quantitative analysis of various means of maintaining inheritance-based diversity on standard domains like the ones we used in this paper. Regarding the multitude of diversity mechanisms present in research, however, it is most important to also point out the results of Wineberg and Oppacher (2003), who most drastically show that “all [notions of diversity] are restatements or slight variants of the basic sum of the distances between all possible pairs of the elements in a system” and suggest that “experiments need not be done to distinguish between the various measures”, a point which we already built upon in our evaluation.

Note that a variety of “meta-measurements” for the analysis of evolutionary processes exist: *Effective fitness* measures the minimum fitness required for an individual to increase in dominance at a given generation (when in competition with the other individuals) (Stephens 1999). It is related to *reproductive fitness*, which is the probability of an individual to successfully produce offspring (Hu and Banzhaf 2010). Both occur at the foundation of productive fitness, but do not include the (computationally overly expensive) diachronical analysis of the overall effect for the end result. Our approach is also comparable to *entropy-based diversity preservation* (Squillero and Tonda 2008), where the positive effect of certain individuals on the population’s entropy is measured and preserved in order to deliberately maintain higher entropy levels. By contrast, our approach is based on the fitness values only (without the need to look into the individuals beyond their genealogical relationships) and thus also cannot be used directly as a secondary goal in evolution but purely as a tool of a posteriori analysis on the effectiveness of other secondary goal definitions.

When we construct the “optimal evolutionary process”, we construct a dynamic optimization problem from a traditionally static one. It is interesting that specifically dynamic or *on-line* (Bredeche et al. 2009) evolutionary algorithms have been shown to benefit from increased diversity especially when facing changes in their fitness functions (Gabor et al. 2018; Grefenstette 1992). While this is obviously intuitive as more options in the population allow for higher coverage of possible changes, the reverse connection (pointed to by this work) is not stated there, i.e., that diversity in static domains may work because even for static domains the optimization process is inherently dynamic to some degree.

Fig. 9 Evolution for the Schwefel problem. Standard evolutionary process using OF shown in black, diversity-aware evolutionary process using AF with Manhattan distance shown in blue. Inherited fitness (purple) and fitness sharing (yellow) shown for comparison. Results averaged over 20 independent runs, the standard deviation is shown in transparent lines



6 Conclusion

We have introduced the novel notion of final productive fitness FPF (and all the definitions it is built upon). We make a theoretical argument that FPF is the goal an optimal evolutionary process should strive for to achieve the best overall results. However, FPF cannot be computed efficiently in advance, producing the need for an approximation. We argue that the well-known technique of augmenting the objective fitness function with an additional diversity goal (when it helps) happens to effectively approximate the theoretically derived FPF (at least better

than just the objective fitness on its own). We have shown this connection empirically on benchmark domains. We argue that this provides first insight into *why* and *how* diversity terms are beneficial to evolutionary processes.

Immediate future work would consist of answering the *when* and *which*: We have tested several domains for evolutionary algorithms and many are too simple to further benefit from explicit diversity-awareness. Maybe FPF can be used to derive a criterion to estimate the usefulness of diversity in advance. Similarly, many mechanism to cater explicitly to diversity exist. While many can be subsumed by the pairwise distance used here (Wineberg and

Oppacher 2003), others may still show different behavior. Their relation to FPF requires further work.

There are means of maintaining diversity without altering the fitness function, most prominently structural techniques like island models (Tomassini 2006; Whitley et al. 1999) or hypermutation phases (Morrison and De Jong 2000; Simões and Costa 2002). As no match for $|AF - FPF|$ can be computed for them, we omitted them in this first analysis. Final productive fitness may be a useful tool to translate these structural means into an effect of the fitness function.

Eventually, there may be even more direct or outright better (compared to using diversity or similarly augmented fitness) approximations of FPF to be found now that we know what we are looking for. The ultimate goal of the research into FPF might be to utilize it directly or indirectly in actually constructing new types of evolutionary algorithms (instead of it “only” helping to explain how well-known types work). We can imagine:

- We could compute FPF for a simplified version of the problem (and the algorithm) for various parametrizations (Eiben and Smith 2003; Mitchell 1998). The FPF’s values could then help evaluate the respective parametrization’s success. However, for most complex problems it is not entirely clear how to derive simpler instances that still preserve the interesting or challenging aspects of their larger counterparts. Furthermore, it is unclear how FPF might provide more information than just using the respective runs’ OF.
- In this paper, we checked how well various established models for fitness functions approximate FPF. Instead, we might now construct new models with the goal to approximate FPF better. Surrogate models have been used in evolutionary algorithms to approximate computationally expensive fitness functions (Gabor and Altmann 2019; Jin 2005; Jin and Sendhoff 2004). In our case, a surrogate model would have to approximate our a-posteriori-approximation of FPF and could then maybe save time for future evaluations. It should be noted that our results make it seem plausible that no general model for FPF on all domains should exist and we cannot train the surrogate as the algorithm goes along since our target metric is only computed a posteriori. However, we might still be able to learn an n -PF surrogate for small n or a similar target metric. In Gabor and Linnhoff-Popien (2020) we already suspect (based on the a posteriori approximation as we do in this paper) that FPF constructs a simpler fitness landscape compared to, quite like surrogates do, suggesting that surrogates may be trained to achieve a similar fitness landscape. That dynamic should be explored in future work.

- In Gabor and Belzner (2017) and Gabor et al. (2018) we introduce *genealogical distance* as a diversity metric for evolutionary algorithms. We show that it provides similar although at times inferior results to distance-based diversity metrics (i.e., the Manhattan and Hamming diversity we use in this paper). However, genealogical diversity may provide means to approximate genealogical relations between individuals making our a-posteriori-approximation much easier to compute (probably at the expense of accuracy). Future work should evaluate if that approach brings any relevant benefits.

Some of these may be applicable to other methods of optimization as well: We suspect that the notion of final productive fitness translates directly to all optimization methods (ranging from simulated annealing or particle swarm optimization to Monte-Carlo tree search and back-propagation) that may or may not already implement means to approximate final productive fitness rather than just objective fitness.

Acknowledgements We thank the anonymous reviewers for their in-depth comments and their contribution to this paper. We also thank Lenz Belzner for intense discussions that helped make this paper possible.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Benchmarks – DEAP 1.3.1 documentation. <http://deap.readthedocs.io/en/master/api/benchmarks.html>. Accessed: 2020-04-30
- Brameier M, Banzhaf W (2002) Explicit control of diversity and effective variation distance in linear genetic programming. In: European Conference on Genetic Programming, pp 37–49. Springer
- Bredeche N, Haasdijk E, Eiben A (2009) On-line, on-board evolution of robot controllers. In: International Conference on Artificial Evolution (Evolution Artificielle), pp 110–121. Springer
- Brown G, Wyatt J, Harris R, Yao X (2005) Diversity creation methods: a survey and categorisation. *Inf Fusion* 6(1):5–20

- Burke EK, Gustafson S, Kendall G (2004) Diversity in genetic programming: an analysis of measures and correlation with fitness. *IEEE Trans Evolut Comput* 8(1):47–62
- Chen JH, Goldberg DE, Ho SY, Sastry K (2002) Fitness inheritance in multi-objective optimization. In: *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*, pp 319–326
- Conti E, Madhavan V, Such FP, Lehman J, Stanley K, Clune J (2018) Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In: *Advances in Neural Information Processing Systems*, pp 5027–5038
- Corno F, Sánchez E, Squillero G (2005) Evolving assembly programs: how games help microprocessor validation. *IEEE Trans Evolut Comput* 9(6):695–706
- Črepinšek M, Liu SH, Mernik M (2013) Exploration and exploitation in evolutionary algorithms: a survey. *ACM Comput Surv* 45(3):35
- Eiben AE, Smith JE et al (2003) *Introduction to evolutionary computing*, vol 53. Springer, Berlin
- Gabor T, Altmann P (2019) Benchmarking surrogate-assisted genetic recommender systems. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp 1568–1575
- Gabor T, Belzner L (2017) Genealogical distance as a diversity estimate in evolutionary algorithms. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp 1572–1577
- Gabor T, Belzner L, Linnhoff-Popien C (2018) Inheritance-based diversity measures for explicit convergence control in evolutionary algorithms. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp 841–848. ACM
- Gabor T, Belzner L, Phan T, Schmid K (2018) Preparing for the unexpected: Diversity improves planning resilience in evolutionary algorithms. In: *2018 IEEE International Conference on Autonomic Computing (ICAC)*, pp 131–140. IEEE
- Gabor T, Linnhoff-Popien C (2020) A formal model for reasoning about the ideal fitness in evolutionary processes. In: T Margaria, B Steffen (eds) *Leveraging Applications of Formal Methods, Verification and Validation: Engineering Principles - 9th International Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes, Greece, October 20–30, 2020, Proceedings, Part II, Lecture Notes in Computer Science*, vol. 12477, pp 473–490. Springer. https://doi.org/10.1007/978-3-030-61470-6_28
- Grefenstette JJ et al (1992) Genetic algorithms for changing environments. *PPSN* 2:137–144
- Hart E, Sim K (2018) On constructing ensembles for combinatorial optimisation. *Evolut Comput* 26(1):67–87
- den Heijer E, Eiben A (2012) Maintaining population diversity in evolutionary art. In: *International Conference on Evolutionary and Biologically Inspired Music and Art*, pp 60–71. Springer
- Hu T, Banzhaf W (2010) Evolvability and speed of evolutionary algorithms in light of recent developments in biology. *J Artif Evolut Appl* 2010:1
- Jin Y (2005) A comprehensive survey of fitness approximation in evolutionary computation. *Soft Comput* 9(1):3–12
- Jin Y, Sendhoff B (2004) Reducing fitness evaluations using clustering techniques and neural network ensembles. In: *Genetic and Evolutionary Computation Conference*, pp 688–699. Springer
- McPhee NF, Hopper NJ (1999) Analysis of genetic diversity through population history. In: *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 2*, pp 1112–1120. Morgan Kaufmann Publishers Inc
- Mitchell M (1998) *An introduction to genetic algorithms*. MIT press, Cambridge
- Morrison RW, De Jong KA (2000) Triggered hypermutation revisited. In: *Proceedings of the 2000 Congress on Evolutionary Computation. CEC00 (Cat. No. 00TH8512)*, vol 2, pp 1025–1032. IEEE
- Morrison RW, De Jong KA (2001) Measurement of population diversity. In: *International Conference on Artificial Evolution (Evolution Artificielle)*, pp 31–41. Springer
- Ng AY, Harada D, Russell S (1999) Policy invariance under reward transformations: theory and application to reward shaping. *ICML* 99:278–287
- Pugh JK, Soros LB, Stanley KO (2016) Quality diversity: a new frontier for evolutionary computation. *Front Robot AI* 3:40
- Rainville D, Fortin FA, Gardner MA, Parizeau M, Gagné C et al. (2012) Deap: A python framework for evolutionary algorithms. In: *Proceedings of the 14th annual conference companion on Genetic and evolutionary computation*, pp 85–92. ACM
- Sareni B, Krahenbuhl L (1998) Fitness sharing and niching methods revisited. *IEEE Trans Evolut Comput* 2(3):97–106
- Simões A, Costa E (2002) Using genetic algorithms to deal with dynamic environments: a comparative study of several approaches based on promoting diversity. *Proceed Genet Evolut Comput Conf GECCO* 2:698–707
- Squillero G, Tonda A (2016) Divergence of character and premature convergence: a survey of methodologies for promoting diversity in evolutionary optimization. *Inf Sci* 329:782–799
- Squillero G, Tonda AP (2008) A novel methodology for diversity preservation in evolutionary algorithms. In: *Proceedings of the 10th annual conference companion on Genetic and evolutionary computation*, pp 2223–2226. ACM
- Stephens CR (1999) “Effective” fitness landscapes for evolutionary systems. In: *Proceedings of the 1999 congress on evolutionary computation-CEC99 (Cat. No. 99TH8406)*, vol 1, pp 703–714. IEEE
- Toffolo A, Benini E (2003) Genetic diversity as an objective in multi-objective evolutionary algorithms. *Evolut Comput* 11(2):151–167
- Tomassini M (2006) *Spatially structured evolutionary algorithms: artificial evolution in space and time*. Springer, Berlin
- Ursem RK (2002) Diversity-guided evolutionary algorithms. In: *International Conference on Parallel Problem Solving from Nature*, pp 462–471. Springer
- Whitley D, Rana S, Heckendorn RB (1999) The island model genetic algorithm: on separability, population size and convergence. *J Comput Inf Technol* 7(1):33–47
- Wineberg M, Oppacher F (2003) The underlying similarity of diversity measures used in evolutionary computation. In: *Genetic and Evolutionary Computation Conference*, pp 1493–1504. Springer

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.