Evolutionary computing to assemble standing genetic diversity and achieve long-term genetic gain

Kira Villiers (ORCID: 0000-0002-6410-3477)[1,*], Kai P. Voss-Fels (ORCID: 0000-0003-0782-366X)[2,1], Eric Dinglasan (ORCID: 0000-0002-4824-5977)[1], Bertus Jacobs[3], Lee Hickey (ORCID: 0000-0001-6909-7101)[1], Ben J. Hayes (ORCID: 0000-0002-5606-3970)[1]

[1] Queensland Alliance for Agriculture and Food Innovation, The University of Queensland, St Lucia, QLD, Australia.

[2] Department of Grapevine Breeding, Hochschule Geisenheim University, Geisenheim, Germany.

[3] LongReach Plant Breeders Management Pty Ltd, Lonsdale, SA, Australia.

[*] To whom correspondence should be addressed.

Contact address: k.villiers@uq.edu.au

Abstract

Loss of genetic diversity in elite crop breeding pools can severely limit long-term genetic gains, and limits gain for new traits that are becoming important as the climate changes, such as heat tolerance. Introgression of specific traits and pre-breeding germplasm is an alternative to introduce new diversity, however, introgression is typically slow and challenging. Here we investigate optimal haplotype selection (OHS), a selection method which retains useful diversity in the population by selecting sets of candidates which between them contain haplotype segments with very high segment breeding values for the target trait. In simulation, the improvements for a yield trait in a closed population beginning with a founder population of wheat genotypes, were compared for OHS, truncation selection on genomic estimated breeding values (GEBVs), and optimal cross selection (OCS). After 100 generations of selection and intercrossing, truncation selection had exhausted the genetic diversity, while considerable diversity remained in the OHS population. In situations where the number of parents crossed each generation was relatively small and the number of progeny per cross was large, gain from OHS ultimately exceeded that from truncation selection. Compared to OCS, selection under OHS maintained more useful diversity, that is diverse segments with high yield local GEBV. Introducing a single cycle of OHS selection early in a recurrent truncation selection program substantially improved long term gains compared with truncation selection. These results indicate that OHS-based parent selection could be a promising method to maintain diverse alleles and achieve higher long term genetic merit in breeding populations.

Keywords: Long-term selection; genomic selection; optimal haplotype stacking; breeding simulation

Key message (30 words or less):

A segment-stacking approach that maintains useful genetic diversity can achieve reasonable rates of gain in small populations and can be introduced for a single generation to produce improvements in genetic value in both short and long-term timeframes.

Introduction

A breeding program can be seen as a targeted reduction in the diversity of a population. That is, undesirable alleles are culled from the pool, and beneficial alleles are increased in frequency. However, a breeding program with no diversity can produce no genetic gains, because every genotype in the population is identical with regards to its allelic constitution and represents the peak of the population's genetic merit. A significant challenge for elite plant breeding programs is to maintain or increase rates of gain over long time periods without exhausting their breeding program's genetic diversity.

Crop breeding programs are increasingly turning to genomic selection (GS) to increase genetic gains. GS involves ranking selection candidates or candidate parents for parent selection on genomic estimated breeding values (GEBVs) calculated from genotypes marker effects estimated from a reference population (Meuwissen et al. 2001). Genomic selection can provide very high rates of gain in some species as it enables early selection of breeding individuals (Voss-Fels et al. 2019). The introduction of GS has boosted genetic gain in both animal and plant breeding programs (García-Ruiz et al. 2016; Guinan et al. 2023; Scott et al. 2021; Wolc et al. 2015; Crossa et al. 2014; Cooper et al. 2014; Bernardo 2021; Hickey et al. 2017).

An emerging strategy to integrate GS in crop breeding programs is rapid recurrent selection on GEBVs, especially in the population improvement component of a two-part breeding program (Gaynor et al. 2017). In a two-part breeding program, rapid breeding cycles in the population improvement component mean good alleles approaching fixation faster, which translates into greater performance gains in the outputs of the other half of the breeding program, the product development pipeline. Technologies key to rapid iteration in the population improvement component are GS, by allowing ranking and crossing of candidates as early as generation F1 without field testing (so that field testing and updating of marker effects can stay in the purview of the separate product development component of the program), and speed breeding (Watson et al. 2018), by shortening the growth period to allow more generations in the same time span. Truncation selection on GEBV is the most commonly used approach to select parents in the population improvement component of breeding programs. While this leads to rapid gains, truncation selection can cause dramatic loss of diversity and often lead to rapid inbreeding in the breeding pool (Lin et al. 2016; Cowling et al. 2019). In dairy cattle, for example, the rate of inbreeding has doubled as a result of the introduction of genomic selection (Scott et al. 2021).

Potentially, the loss of diversity caused by selecting this way could be mitigated by introgressing genetically distant germplasm (Gorjanc et al. 2016, Allier et al. 2020, Sanchez et al. 2023). However, truncation selection is not well suited to introgression scenarios. Diverse or landrace germplasm often lacks important agronomic traits compared to elite varieties that have been developed through generations of targeted selection. The diverse germplasm's mediocre GEBVs for key traits would result in these individuals being culled from the breeding program, even though they may carry beneficial alleles not yet present in elite varieties. The benefits of the diverse germplasm are 'masked' by the whole-genome GEBV scores (Gorjanc et al. 2016). A whole-genome perspective loses the locus-specific information that could distinguish the germplasm's usefulness based on individual alleles, and truncation selection does not take any diversity measures into account.

Optimal contribution selection (OCS) is a strategy for balancing genetic gain and genetic diversity in animal and plant breeding programs (Wray and Goddard 1994; Meuwissen 1997; Woolliams et al. 2015; Lin et al. 2017; Gorjanc et al. 2018). In OCS, breeders balance the minimisation of coancestry, a measure of loss of genetic diversity, against the maximisation

of genetic gain, for example improvements in the trait GEBV. The possible weightings of these two objectives produce a Pareto frontier of OCS strategies for different goals or priorities. OCS programs determine genetic contribution proportions from each potential parent, or can be adapted to provide mate allocations as “optimal cross selection” (Kinghorn 2011). OCS was designed for animal breeding, but is also seeing increasing use in plant breeding. Different plant breeding program designs based on OCS have been proposed (Cowling et al. 2017; Cowling et al. 2019), and it has been applied in forestry (Kerr et al. 2015), common bean breeding (Saradadevi et al. 2021), and spring canola breeding (Cowling et al. 2023) projects.

To take advantage of genetic marker data, OCS can be carried out using GEBVs as the measure of genetic value and the genomic relationship matrix (GRM) as the measure of relatedness. One variation on OCS is Tiret et al. (2021)’s genomic OCS (GOCS), which adds a term to heterozygote x heterozygote crosses in the GRM and in doing so, sees improvements over OCS in genetic gain at the same coancestry measure when using non-inbred populations. Because OCS for selecting mating pairs is often used to produce full sibling families out of which the best are then selected, Akdemir & Sánchez (2016) developed their Genomic Mating variant of OCS, which replaces the genetic value term in the constraints with a “risk” term composed of a pair’s GEBVs and expected progeny variance. A similar strategy is Allier et al. (2019)’s usefulness criterion parental contribution (UCPC) method, which uses OCS-like parental contribution constraints but defines the genetic value objective using expected progeny mean and variation. Both GOCS, over 20 generations of recurrent selection (Tiret et al. 2021), Genomic Mating, over 20 cycles of recurrent selection (Akdemir & Sánchez 2016), and UCPC, over 60 generations of recurrent selection (Allier et al. 2019, Sanchez et al. 2023), slightly outperformed OCS for genetic gain. Recurrent OCS itself can provide rates of gain in line with or outperforming truncation selection within 15-20 generations of simulation, after a slight lag for the first few generations (Akdemir & Sánchez 2016, Sanchez et al. 2023).

While GS and OCS use whole-genome ratings to guide selection, selection methods could also be designed around the concept of collecting beneficial alleles into one genotype. Stacking, “gene pyramiding” or “genotype building” approaches in their direct form have not been widely applied in plant breeding (Han et al. 2017). In stacking approaches, genetic gain relies on beneficial recombination events causing the accumulation of good alleles or good segments in the same genotype (Dekkers and Hospital 2002; Servin et al. 2004). The minimum population sizes to guarantee fixation of desirable alleles or segments grow exponentially with the number of alleles or segments being targeted (Hospital 2003; Servin et al. 2004). Therefore, stacking approaches do not seem suited to use together with GS, which uses thousands of markers. However, by bundling markers into chromosome segments, and selecting on chromosome segment haplotypes, this kind of selection could become feasible, as demonstrated by Kemper et al. (2012). This approach assumes every chromosome segment contains at least one mutation affecting the trait, and is effectively a version of gene-pyramiding adapted to a quasi infinitesimal model of quantitative trait variation (many mutations of small effect spread across the genome). The optimal haplotype stacking (OHS) approach described by Kemper et al. (2012) was to selecting a set of candidates whose pool of haplotype segments could be stacked into the best “ultimate genotype”, the highest-scoring genotype possible out of the segments in the population.

A number of authors have extended this concept for plant breeding applications. The optimal haploid value (OHV) (Daetwyler et al. 2015) of a candidate is the GEBV of the best doubled haploid that could be produced from that candidate out of its diploid marker block haplotypes. Selection on the OHV of candidates in a doubled-haploid-based breeding

program outperformed selection on candidate's GEBV for both genetic gain and diversity retention (Daetwyler et al. 2015). Variants of OHV selection exist that extend the method by, instead of calculating a score for each candidate, calculate a score for each potential mating pair or each potential selection of candidates. Han et al. (2017) treated all alleles as either desirable or undesirable in order to define the predicted cross value (PCV) as the probability that the progeny of a cross between two candidates carries only desirable alleles. In a simulated introgression scenario, selecting pairs to cross on the basis of their PCV allowed the alleles to be introgressed into the population faster than truncation selection on GEBVs (Han et al. 2017).

The OHV and PCV strategies are based on the chromosome segments carried by individual candidates. At the population level, Goiffon et al. (2017) compared Kemper et al. (2012)'s method (OHS) to other selection methods in simulation, including their own derived method OPV. Selection on the optimal population value (OPV) of Goiffon et al. (2017) is the optimisation task of selecting a set of complementary candidates whose pool of haplotype segments can be stacked to produce the doubled haploid with the highest GEBV. It could be seen as a version of OHV where the whole population's haplotype segments are available to combine into the best doubled haploid, rather than only the haplotype segments of an individual candidate. Over most replications of Goiffon et al. (2017)'s simulation experiment, selection on OPV produced the largest genetic improvement after 10 generations of recurrent selection, followed by OHS. Both outperformed selection on individual-candidate OHV and GEBV. They did not, however, test the performance of OHS in recurrent applications or over long time frames. OHS and OPV can be considered an alternate approach to stacking based on evolutionary computing solvers, rather than on an exact solver from an operations research approach like Han et al. (2017).

Here, we investigate the performance of OHS, as implemented in Kemper et al. (2012), when applied over successive generations in a closed population, and across a wide variety of parameters including number of parents available for crossing and family sizes. The aim was to provide breeders some insight into when OHS might be useful strategy to accelerate genetic gains and maintain genetic diversity. Recurrent rounds of OHS selection were performed in simulation, starting with real wheat founder genotypes. Performance of OHS was compared to recurrent rounds of truncation selection and OCS, and in some cases OPV, in matched simulations. Performance was measured in terms of achieving genetic gain, maintaining genetic diversity, and preserving high-quality haplotype segments. We hypothesised that OHS would outperform the other two methods in maintaining useful diversity, and therefore, in a closed population, over long time horizons (to allow for stacking to occur), could produce higher genetic gain. We also hypothesised that using OHS to select founder lines would provide good foundation of useful standing variation for a breeding program.

Methods

Base population

LongReach Plant Breeders provided phased genotypes of a set of 16,565 diverse elite wheat germplasm lines. The genotypes spanned 5112 SNPs on 21 chromosomes (with a total genome length of 2994 centimorgans), were provided by breeding company LongReach Plant Breeders. A set of additive effect effects for the 5112 SNPs, describing grain yield, were also provided. Those SNP yield effects were from best linear unbiased predictions (BLUPs) on data from 186 field trials between 2012 and 2018. For the simulations in this paper, these were used as true SNP effect values that stayed constant

over all generations. All GEBVs discussed in this study are based on this set of marker effects. Hence, the 'GEBVs' in the simulations in this study function as true breeding values.

Two base populations of 50 lines were prepared from the diverse germplasm. One set (GS-selected founders, abbreviated GS_f) was selected by truncation selection – i.e., the 50 genotypes with the highest GEBVs. The other set (abbreviated OHS_f) were selected using OHS on the genotyped population. The genomic relationships among the base population genotypes are shown in Figure 1. Matching simulation experiments were carried out on both base populations, and the differences in the results based on the base population are discussed.

Optimal Haplotype Selection

If a GEBV is a whole-genome score representing the estimated contribution of QTLs near to a genetic marker to the phenotype of a candidate organism, we can define a corresponding "local GEBV" that is the estimated contribution of QTLs in a particular genomic region of one haplotype of that candidate:

$$\text{local GEBV} = \sum_{k \in m} h_k e_{1k} + (1 - h_k) e_{0k}$$

where m is the set of SNP markers belonging to the segment/marker block for which the local GEBV is calculated, h_k is the allele of one marker (either 0 or 1), and e_{0k} and e_{1k} are the additive marker effects of alleles 0 and 1, respectively, on marker k, provided by LongReach Plant Breeders.

We call the genome region on which a local GEBV is calculated a haplotype block. For a diploid genome, there are two haplotype blocks and so two local GEBVs for any specified genomic region. Given a set of marker blocks that partition the genome and assuming additive gene action only, the sum of the local GEBVs for those marker blocks add up to the candidate's GEBV.

The goal of OHS is to maximise the score of the "ultimate genotype" that can be constructed out of haplotype blocks found in its input candidate set. The ultimate genotype of a subset is the one that has the highest GEBV, so the one that has the two genetic sequences at each marker block that produce the two highest local GEBVs for that marker block, with the additional restriction that those two haplotype blocks should originate in separate members of the candidate subset.

To select using OHS, a user chooses the number n of candidates they wish to select out of their candidate pool, and chooses a set of marker blocks for which they can provide a pair of local GEBVs for each block and for every member of the candidate pool. The method will advise the user which n candidates should be selected to create a sample with the highest stacking potential (highest-scoring ultimate genotype). Formally, the function that OHS optimises is:

$$\text{Maximise} \sum_{s \in S} \sum_{b \in B} a_{sb}\, x_{sb}, \ \text{ subject to } \sum_{s \in S} a_{sb} = 2 \ \text{ for all } b \in B$$

where S is any n-member subset of the candidate pool; B is the set of marker blocks on which the potential founders are scored; x_{sb} is the higher of the candidate s's two local GEBVs for marker block b; and a_{sb} are a collection of binary (0 or 1) decision variables that are constrained to select exactly two haplotype blocks from S for each marker block b, out of which the ultimate genotype of the candidate subset S is constructed.

The fact that the x_{sb} term is the larger of a candidate's two haplotype's local GEBVs for a block ensures the two haplotype blocks that stack to make the optimal genotype's information for that block must come from two different candidates. This does not require that the two haplotype blocks are distinct and the optimal genotype is heterozygous, but does encourage that the choice of more candidates should capture more diversity. If there was one very elite candidate, and the fitness function could choose that the optimal genotype carried all of the haplotype blocks of that elite candidate in both haplotypes, then the rest of the selected candidates would be random choice. With the restriction on the optimal genotype descending from two candidates at each block, the elite candidate will still be selected but the remainder of the selection will still have value (be carrying good, though not quite as good, haplotype blocks).

It might seem that in plant breeding where inbred or doubled haploid lines are the goal, the two-haplotype-per-chromosome segment constraint is unnecessary. Goiffon et al. (2017)'s OPV function is a modification of the OHS/Kemper et al. (2012) fitness function to remove the restriction, meaning it selects a set of candidates for the best possible fully-homozygous stack. In the same scenario as described above, with a particularly elite candidate, the rest of the candidates selected by OPV would also be random choice. Comparisons between the performance of OPV and OHS will be discussed later in this paper.

Choosing a subset that contains the segments to produce the highest-scoring ultimate genotype is a highly combinatorial problem. The number of possible subsets grows rapidly with the size of the population and the size of the desired subset, such that it is completely infeasible to explore all possible subsets to find the optimum. Even the task of choosing 50 founders from a base breeding population of 200 has over 10^{47} possible subsets of 50. Finding the global maximum ultimate genotype GEBV is therefore too large of a task to solve analytically. Kemper et al. (2012) used an evolutionary search strategy to produce a 'good enough' solution in a defined number of iterations. The implementation of OHS in the genetic algorithm (GA) that they used to select founders in their study is the same as the implementation used here. Further information on the GA implementation's design is available in Kemper et al. (2012)'s Supplementaries.

Preliminary trials showed that running Kemper et al. (2012)'s GA with an internal population size of 1000 chromosomes (candidate subsets S) for 2500 generations of iteration reliably led to acceptable convergence of the fitness function. These hyperparameter values were used when collecting all remaining OHS results in this paper.

The number of marker blocks used in OHS simulations was fixed at 105 (segments containing all the markers in each 1/5th of each chromosome by length). Preliminary trials confirmed that the number of marker blocks does affect the performance of OHS: this is discussed in the discussion.

Simulation Design

Design of the simulations in this investigation did not aim to mimic the design of a real-world breeding program, rather they were designed to give insights into the situations in which OHS might be a useful strategy. Simulations involved recurrent selection over a long (100 non-overlapping generations) time frame, to test the long-term limits of the selection strategy. The population at each generation were F1 crosses of the previous generation; that is, no selfing or haplotype-doubling generations occurred in simulation, and candidates were not inbreds.

The tool genomicSimulation (Villiers et al. 2022) was used to simulate the genotypes of offspring of crosses, as well as to calculate the GEBVs and local GEBVs on which selection

occurred. genomicSimulation assumes crossovers in meiosis occur with uniform probability across the length of each chromosome, and does not simulate mutation. It was necessary that the simulation tool track the haplotype phase of the genotype, because calculating local GEBVs for OHS requires this information; it is necessary to know which alleles are in which haplotype and so which alleles contribute to which of the two local GEBVs for a marker block and a candidate. This is especially relevant in a population of highly heterozygous candidates, like the recurrent selection without selfing produces in this simulation, where the two alleles for a marker are often different.

A population size n (n = 2, 10, 25, 50 or 100) was maintained at every generation (Figure 2). To progress to the next generation, the population had to be replaced with a population made up of its progeny, simulated based on the parent genotypes and selected from according to one of the selection methods tested. Genotypes of offspring from the half diallel of $n(n-1)/2$ crosses between pairs of candidates were simulated. Different simulations tested different numbers of progeny produced from each crossed pair in the half diallel (p = 1, 10 or 100). genomicSimulation was then used to produce the table of local GEBVs of the $pn(n-1)/2$ candidates in the offspring pool. The local GEBVs were based on the 105 chromosome segments containing the markers within each 1/5th of the length of each chromosome, as previously described. Novel haplotype blocks and local GEBVs could appear between generations, even though marker effects and chromosome segments blocks were fixed, because genomicSimulation's meiosis simulation allowed recombination within blocks as well as between blocks. The table of local GEBVs of the offspring pool and the request to select n of those candidates were then given to Kemper et al. (2012)'s GA, which implements OHS. The n candidates from the offspring pool selected by OHS then became the parents of the next generation.

Three other selection methods were simulated to compare to OHS. In simulations of truncation selection (Figure 2a), the same simulation of the genotypes of the half-diallel of crosses of the previous generation's selected candidates would be performed, creating an offspring pool of $pn(n-1)/2$ candidates. genomicSimulation was then used to calculate the GEBVs of those candidates, based on the 'true' additive marker effect set, and the n candidates with the highest GEBV were selected to become the parents of the next generation.

At each generation in the simulations of OCS (Figure 2c), the genomic relationship matrix (GRM) of the n candidates was generated and their GEBVs calculated by genomicSimulation. The tool AlphaMate (Gorjanc and Hickey 2018) was used to select optimal matings from those n parent-candidates. genomicSimulation was then used to simulate the genotypes of the offspring resulting from the mating plan produced by AlphaMate. In choosing the best matings, an OCS implementation considers the possible diallel of crosses across two criteria: a measure of coancestry and a measure of genetic gain. The method allows for different weightings of the relative importance of increased genetic gain versus reducing coancestry, depending on the requirements of the particular situation. In AlphaMate, this weighting is represented as the angle on the Pareto frontier, where 0° is ignoring genetic gain entirely in favour of minimising coancestry, and 90° produces a mating plan which ignores coancestry in favour of increasing genetic gain.

Rather than analyse all possible weightings of OCS targets, three OCS targets were defined that would allow for useful comparisons: an OCS weighting angle whose initial rate of genetic gain best matched the initial rate of gain of truncation selection; an OCS weighting angle whose initial rate of genetic gain best matched the initial rate of gain of OHS; and an OCS weighting angle whose initial rate of coancestry increase best matched the initial

coancestry increase of OHS. The “initial rate of increase” to which the OCS angles were matched was the rate of increase in the first 5 generations of selection. To find which angles on the Pareto frontier corresponded to these three targets for each population size n, a grid search of OCS target angles with a coarseness of 5 degrees was performed, followed by binary search to a coarseness of 0.5 degrees. This gave the OCS target angles to 0.5 degree accuracy. The located angles, alongside the factorial design to explore parameters n and p in OHS and truncation selection simulations, can be found in Table 1.

Because there are n pairs in the OCS crossing plan and it is the n offspring of those pairs that become the parents of the next generation (Figure 2), the progeny per cross parameter must be 1 for OCS simulations or else the population size would be larger than declared. Initial trials showed that AlphaMate would often select the same parent, or the same pair of parents, multiple times over in its crossing plan, especially for target angles that highly favoured genetic gain. Clearly the contribution from that parent, or pair of parents, was favoured. Therefore, we consider that the OCS method is capable of choosing multiple (indeed, of choosing an unlimited number) of progeny per pair in AlphaMate. For the target angle grid search, the OCS condition was compared to the OHS or truncation selection condition with the same n, and $p = 100$, or 10 if the 100 progeny per cross condition does not exist for that n (Table 1).

Note that we performed OCS at the parent level, rather that selecting out of offspring pool like with OHS. This was chosen because there was no obvious interpretation for the case where, under optimal contributions, more than one-part contribution was requested from a single candidate, because candidates are immediately used as parents of the next generation. There is no non-crossing/selfing generations that could be used for bulking the descendants of a particular candidate.

Finally, simulations were run that used the same design as OHS simulations (Figure 2b) but replacing the OHS fitness function (selecting two haplotypes per chromosome segment) with a near-identical fitness function that only selected one haplotype per chromosome segment (as in Goiffon et al. (2017)’s OPV).

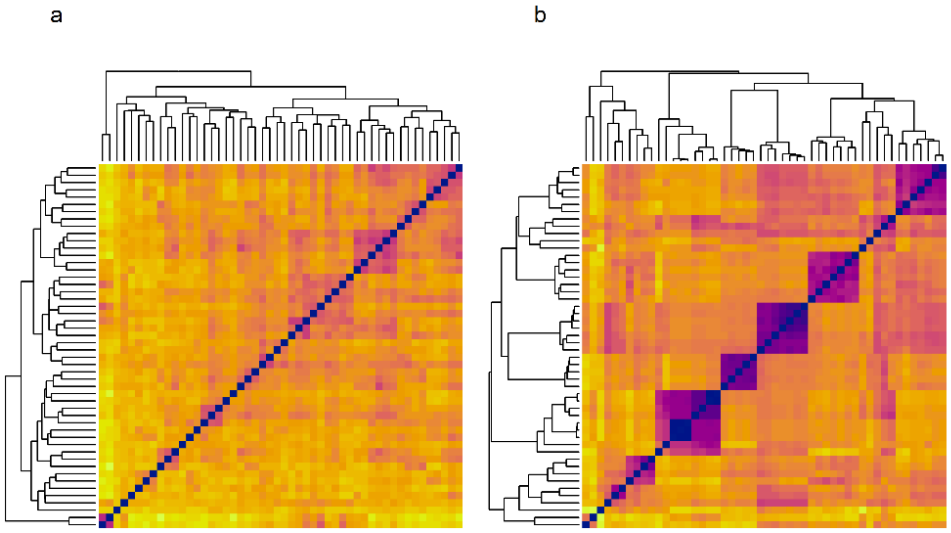

Fig. 1: Summary of the base population wheat genotypes. (a) Heatmap of Roger’s genetic distance between the genotypes of the OHS-selected founding population (that is, the 50 candidates selected by Kemper et al. (2012)’s genetic algorithm (GA) from a large diverse population). Dark colours indicate a low genetic distance

between pairs of candidates. (b) Corresponding heatmap for genetic distance between GS-selected founding genotypes (that is, the 50 candidates selected by truncation selection from the large diverse population).

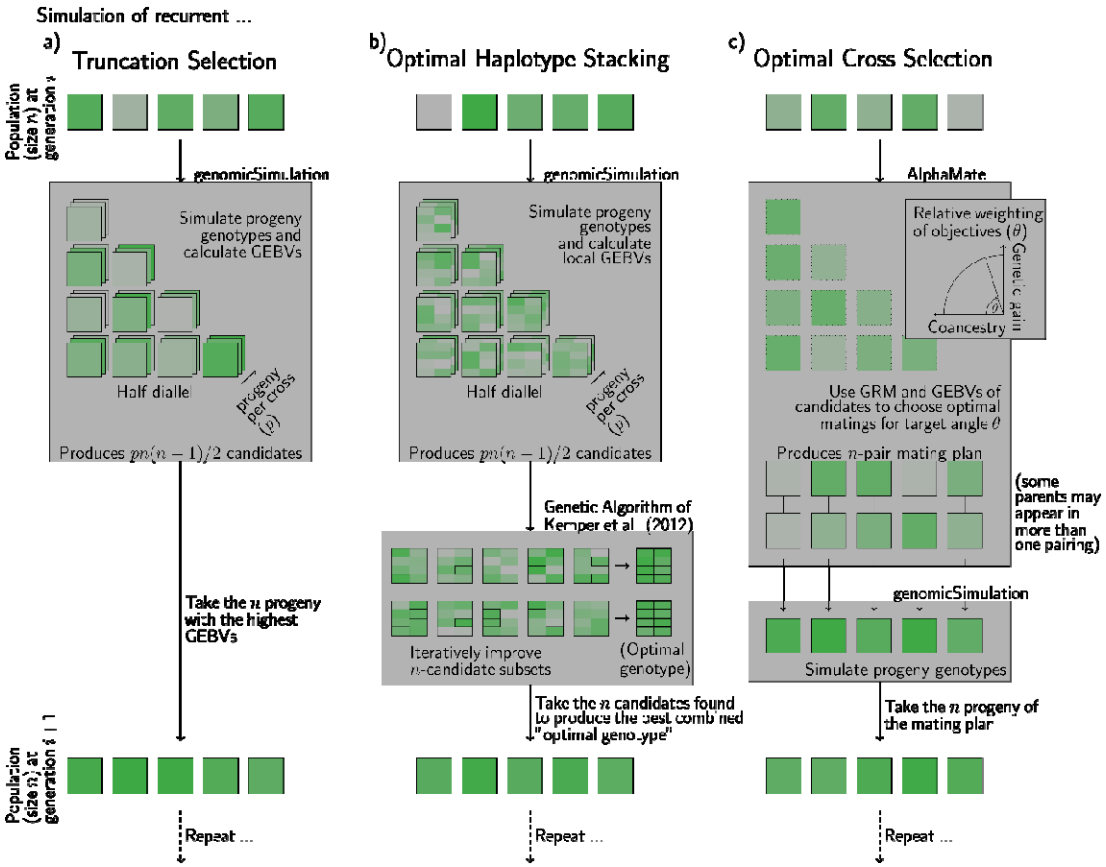

Fig. 2: Structure for the three compared simulation experiments. (a) The recurrent truncation selection simulations involve simulating the genotypes and GEBVs of full-sibling progeny of each possible cross between pairs of candidate genotypes in genomicSimulation (Villiers et al. 2022) , followed by selecting the with the highest GEBVs from the resulting pool of offspring and installing those candidates as the founders of the next generation. (b) OHS (optimal haplotype stacking) simulation structure is the same as the truncation selection procedure for the first step of simulating full-sibling families for each pair in the half-diallel, but instead of selecting on whole-genome GEBVs, selection is done using Kemper et al. (2012)'s genetic algorithm (GA) based on local GEBVs to find a set of candidates in the offspring pool that contain haplotype blocks that can stack to the best "optimal genotype"(c) OCS (optimal cross selection) simulation structure involves using AlphaMate (Gorjanc and Hickey 2018) to choose mate allocations from the set of founder genotypes that would best suit a strategy balancing genetic gain and efforts to minimise coancestry according to a target angle θ between the two objectives. The next generation is then produced by simulating the genotype of one offspring from each of the crosses in the mating plan in genomicSimulation. Across all selection conditions, the selected candidates of each selection cycle become the parents of the next generation, so that simulation tests the performance of the9ifferent methods over long-term recurrent selection.

Table 1: Factorial design of simulation experiments. Each optimal haplotype stacking (OHS) or truncation selection condition can be simulated with a population size of 2, 10, 25, 50 or 100, and 1, 10, or 100 full-sibling progeny per cross ("ppc") in the half-diallel. Conditions with population size 2, and 1 progeny per cross, were not simulated: there is only one possible pairing in the half-diallel of a population with 2 members, so producing only one progeny per cross at that population size is below the replacement rate needed for the next generation. The output format of the genetic algorithm implementation of OHS (Kemper et al. (2012)) was incompatible with candidate pools of more than 99 999 members, so OHS conditions with population size 50 with 100 progeny per cross and population size 100 with 100 progeny per cross were not simulated. Preliminary trials of the optimal cross selection (OCS) software, AlphaMate (Gorjanc and Hickey 2018), showed that AlphaMate would often use a pair or parent multiple times in its mating plan. It was therefore decided to let AlphaMate choose pairs to cross for a population size , and allow it to choose how many of those should be progeny of the same cross.

OCS uses a target angle parameter to define how to weight the two objectives of increasing genetic gain and minimising coancestry gain. A few conceptual targets (matching the initial rates of genetic or coancestry gain of the other selection conditions) were chosen for OCS, such that it would be easy to see the deviations in behaviour between OCS and the other selection conditions. Grid search with a coarseness of 5 degrees, followed by binary search to a coarseness of 0.5 degrees, was used to identify the appropriate target degrees for each conceptual target. The identified target angles are shown in the table.

		Select 2	Select 10	Select 25	Select 50	Select 100
OHS	1 ppc		X	X	X	X
	10 ppc	X	X	X	X	X
	100 ppc	X	X	X		
Truncation	1 ppc		X	X	X	X
	10 ppc	X	X	X	X	X
	100 ppc	X	X	X	X	X
OCS	Match truncation initial coancestry gain		10°	30°	30°	35°
	Match OHS initial GEBV gain		45°	75°	80°	86°
	Match OHS initial coancestry gain		70°	78°	82.5°	86°

Summary and analysis

Simulation of each condition was run independently 20 times for all but the largest OHS scenarios (100 progeny per cross, and a population size of 50 or 100), which, due to runtime issues, were only replicated 10 times. Four summary statistics were calculated for the population at each generation of each simulation trial: mean GEBV, upper selection limit (the GEBV of the ultimate candidate that could be constructed at the current generation, Cole & VanRaden 2011), mean coancestry, and a gene diversity index from Nei (1973). The results show these summary measures averaged across replications.

The first summary statistic, the mean GEBV, is the mean of the breeding values of all the genotypes in the population at a given generation.

The second is Cole & VanRaden (2011)'s upper selection limit: the GEBV of the genotype created by taking the best allele present in the population at every marker. Since the marker effects in this study are purely additive, that genotype will be fully homozygous. Assuming no introduction of fresh diversity, this is the maximum GEBV that descendants of this population could achieve. Unlike the stack-of-segments ultimate genotype optimised in OHS, the upper selection limit is based on the best possible choices at the marker level, so may include chromosome segments that do not exist yet.

The third summary statistic, the mean coancestry of the population at a given generation, is calculated as twice the mean of the values in the genomic relationship matrix, including the diagonal, with the reference population (base) being the founder set, GA_f or GS_f (Gorjanc and Hickey 2018; Isik, Holland, and Maltecca 2017).

Finally, the last statistic was the mean of Nei (1973)'s gene diversity index H_e among a given generation's population, calculated as follows:

$$\frac{1}{|C|}\sum_{g\in C} H_e = \frac{1}{|C|}\sum_{g\in C}\left(1-\sum_{\mathrm{i}=1}^{\mathrm{k}} \mathrm{p}_{\mathrm{i}}^2\right) = \frac{1}{|C|}\sum_{g\in C} 2\mathrm{p}_{\mathrm{A}}\mathrm{p}_{\mathrm{T}}$$

where g is a genotype in C, the population of selected genotypes of a given generation, and p_{A} and p_{T} are the proportions of the alleles of that candidate that are (respectively) A and T.

The simplification is possible thanks to all markers being SNPs (thus $k = 2$), and the assumption of Hardy–Weinberg equilibrium.

Results

Truncation selection versus OHS for genetic gain and maintaining genetic diversity

The base population selected by OHS was considerably more diverse than the base population selected by truncation on GEBVs (Figures 1a and 1b). OHS selection samples more evenly from the available genetic diversity. Conversely, truncation selection strongly favours offspring of only a few founders. This effect is more obvious when the population has more diversity (Figure 3). Compare the more uniform proportion of first-generation offspring of each founder chosen under OHS selection, to the favouring of the offspring of a few founders under truncation selection in Figure 3a.Figure 4 compares the average GEBV of the population at each generation under the aforementioned four simulation conditions: truncation-selected base population and recurrent OHS selection; truncation-selected base population and recurrent truncation selection; OHS-selected base population and recurrent OHS selection; and OHS-selected base population and recurrent truncation selection. Selecting the initial founder population using OHS gives a lower first-generation GEBV but produces higher GEBV scores in the long-term, because rates of gain are steeper under either recurrent simulation strategy. Recurrent truncation selection on the more diverse base population (selected by OHS) surpasses the mean GEBV of truncation selection on the less-diverse starting population (selected by truncation selection) after 12 generations. Recurrent OHS selection produces higher long-term results than recurrent truncation selection. The OHS approach begins to outperform the truncation strategy on the same base population after 70 generations, for the set of parameters displayed: 10 progeny per cross, and selecting 25 genotypes each generation.

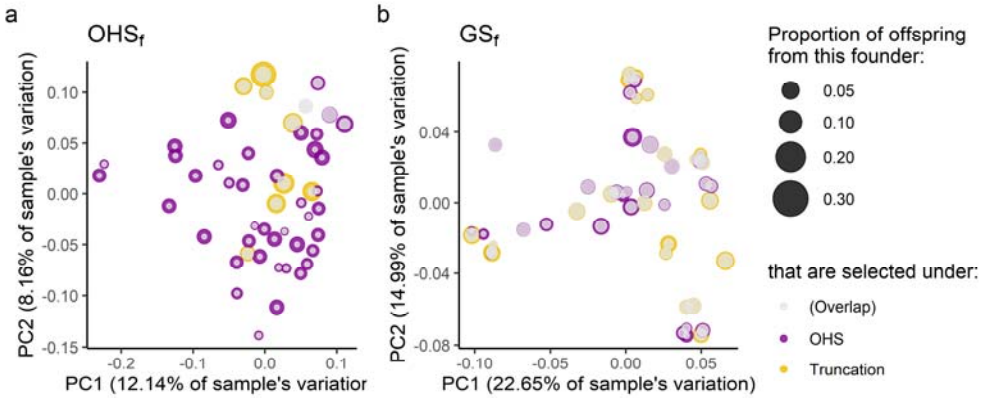

Fig. 3: Diversity of founder population and proportion of lines selected by truncation and optimal haplotype stacking (OHS) selection. (a) Principal component plot of OHS_f, the founding population of 50 lines selected by OHS from a large diverse population. The size of the scatter points corresponds to the proportion of the selected progeny that have that founder as a parent. Proportions of first generation offspring selected by OHS and of first generation offspring selected by truncation selection are shown in different colours. (b) Corresponding principal component plot of GS_f, the founding population of 50 lines selected by truncation selection from the large diverse population.

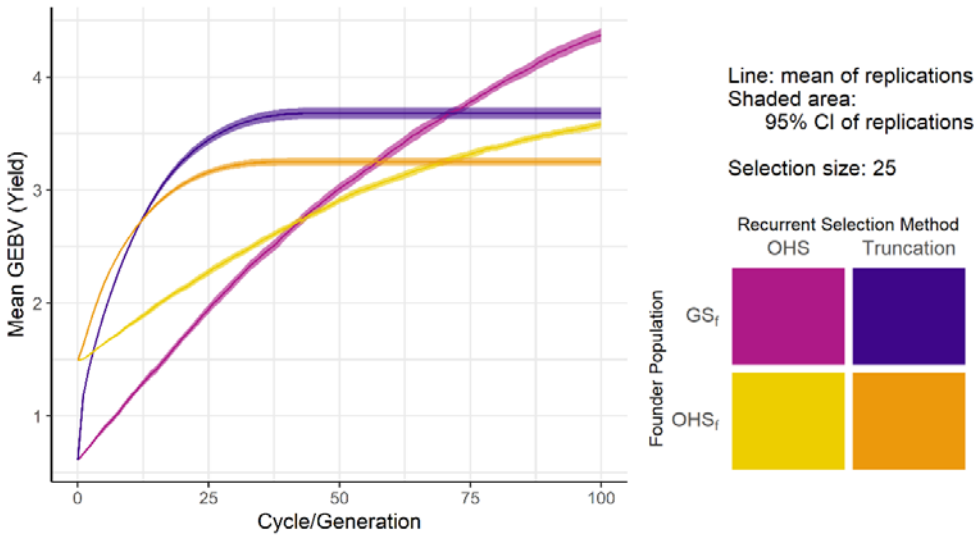

Fig. 4: Trends results of recurrent truncation and OHS selection in simulation, for different founder population diversity. Plot shows each generation's mean and 95% confidence interval of GEBV, comparing simulations starting with OHS-selected (OHS_f) and truncation-selected (GS_f) founder sets and running 100 generations of OHS recurrent selection or recurrent truncation selection. The simulation runs in this plot use the parameters 10 progeny per cross and recurrent selection size of 25.

Figure 5a) is a plot of trends in genetic gain over the 100 generations, comparing OHS and truncation selection over different numbers of progeny per cross. Figure 5b) shows the corresponding trends in genetic diversity, and Figure 5c) the trends in upper selection limit. Figure 5a) shows the same general profiles as Figure 4: while recurrent truncation selection rounds result in very high initial rates of gain, the yield GEBV soon reaches a plateau and sees no improvement thereafter. 5b) demonstrates that the plateau occurs when all the diversity of the population is exhausted, so no further improvement is possible. Meanwhile, OHS selection conditions show substantially slower initial rates of gain in Figure 5a) but also significantly lower rates of loss of diversity in Figure 5b), such that at the end of 100 generations populations had higher GEBVs and a retained stock of diversity they can use to continue improving into the future. From Figure 5c), it is clear that the initial high rates of gain of truncation selection cause the loss of much useful diversity (seen in falling upper selection limits). OHS maintains more genetic diversity over a longer time period.

To investigate the potential practical application of a single generation of OHS alongside truncation selection, Figure 5 also includes a 'hybrid' selection condition, in which the first generation is selected with OHS, but following generations are selected by truncation. Figure 4 and 5a show that even a single generation of OHS selection instead of truncation selection in the early stages of a program, whether to select germplasm out of a larger pool, or to select out of a pool of cross progeny, improves long-term genetic gains. The 'hybrid' selection results show a dramatically higher GEBV plateau, and the plateau is achieved no more slowly than other truncation conditions. The effect is decreased if the OHS generation is in later generations in these simulations, suggesting the benefit depends on diversity in the population. The 'hybrid' selection condition lags slightly behind the corresponding truncation selection condition for the first generation, but surpasses the corresponding truncation selection condition at around generation 10. The 'hybrid' selection condition therefore surpasses pure recurrent truncation selection in both the long-term and shorter term.

Figures 6) and 7) provide summaries of the same three summary statistics as Figure 5) at specific points in time in the simulation runs, across the different combinations of selection size and progeny per cross parameters. Figure 6) shows the summary after 10 generations of selection, and Figure 7) the summary after 100 generations. OHS selection benefits from more progeny per cross, but while, in early generations, truncation selection conditions with more progeny per cross likewise show higher GEBVs, in later generations, after the plateau is reached, truncation selection conditions with more progeny per cross have lower GEBVs than their few-progeny counterparts. While truncation selection benefits from higher population sizes for achieving higher rates of genetic gain, OHS selection benefits, down to a point, from smaller population sizes. When the selection size is 2, both selection conditions exhaust their diversity too fast to see much genetic gain. But GEBVs from OHS selection are higher when the size of the selection is 10 than when it is 25, and higher when it is 25 than when it is 50 or 100. Larger selection sizes, however, allow OHS to maintain more diversity than smaller ones. By generation 100, all genetic diversity with truncation selection was exhausted, but all OHS conditions with population sizes greater than 2 maintain diversity under Nei (1973)'s index, and had higher upper selection limits, suggesting they had maintained useful diversity to use to continue to build improvements.

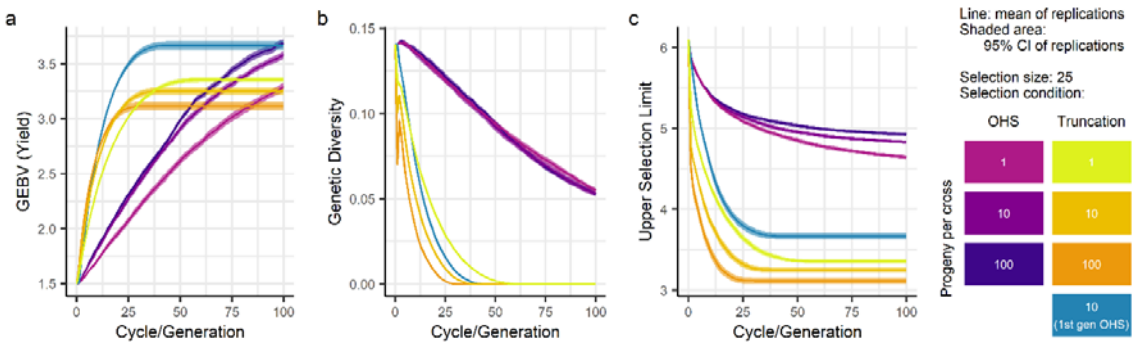

Fig. 5: Comparison of truncation selection and OHS in recurrent selection in a closed population over three performance criteria. (a) Average GEBV of the population over generations. (b) Genetic diversity over generations (Nei (1973)'s expected heterozygosity measure). (c) Upper selection limit (Cole & VanRaden 2011) at each generation. Results are from simulation with a recurrent selection size of 25. Represented in the plots are the following conditions: selection by OHS with 1, 10 or 100 progeny per cross; truncation selection with 1, 10 or 100 progeny per cross; and a 'hybrid' condition with 10 progeny per cross, where selection is done by OHS in the first generation and by truncation for every generation thereafter.

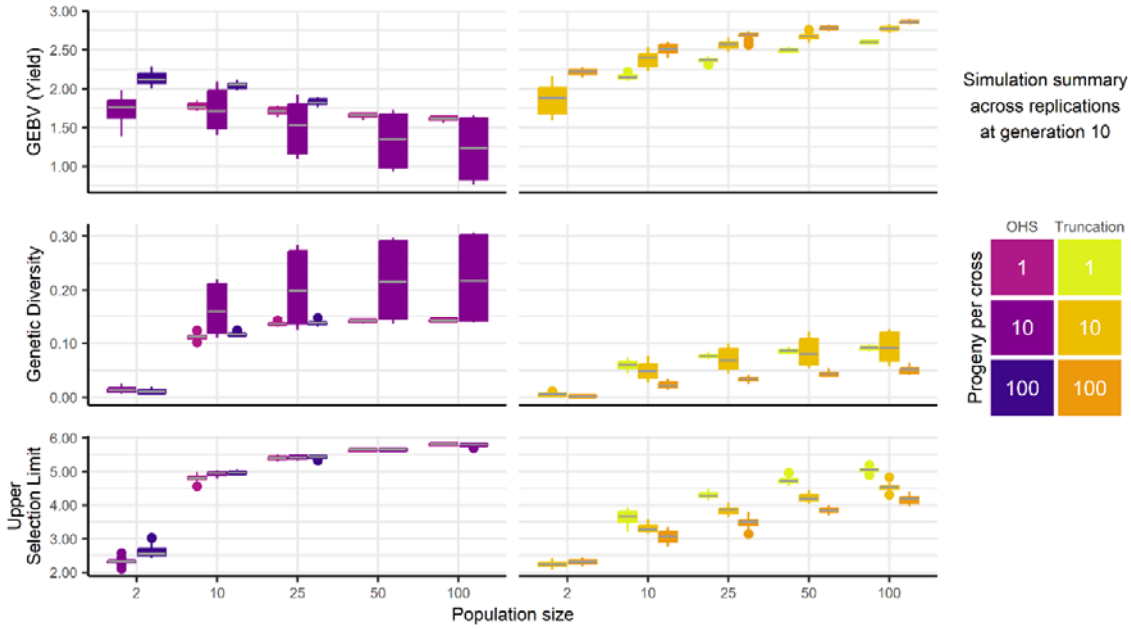

Fig. 6: Comparison of optimal haplotype stacking (OHS) and truncation selection after 10 generations, across all combinations of population size and progeny per cross parameters in the factorial design.

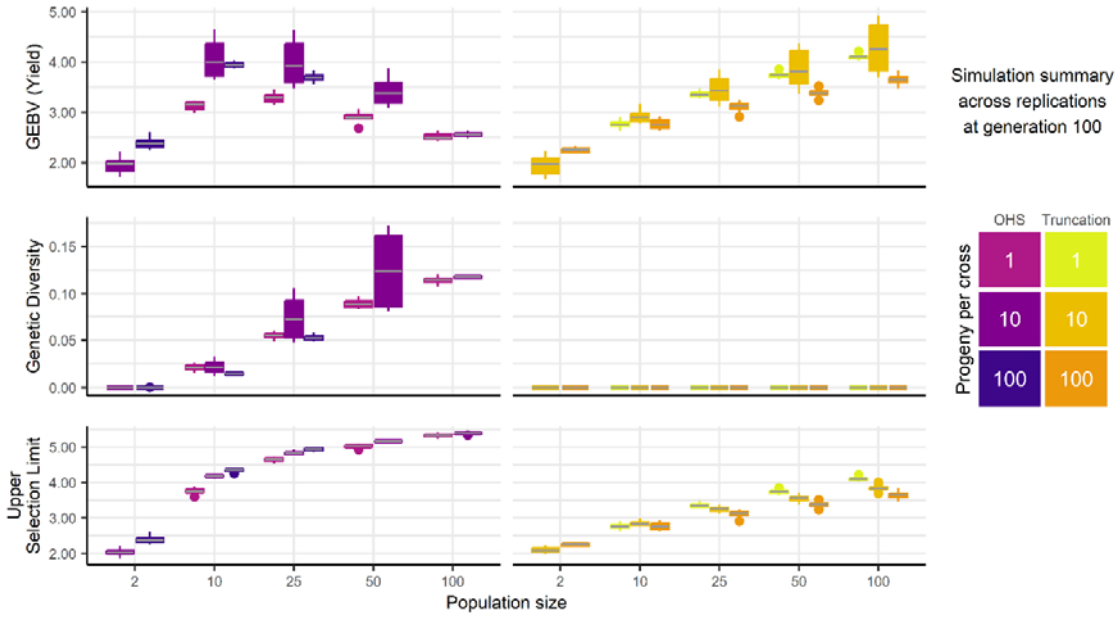

Fig. 7: Comparison of optimal haplotype stacking (OHS) and truncation selection after 100 generations, across all combinations of population size and progeny per cross in the factorial design. For all truncation selection scenarios, genetic diversity was exhausted (0) by generation 100.

Optimum cross selection versus OHS for genetic gain and maintaining genetic diversity

The OCS condition matched to truncation selection's initial rate of coancestry increase shows the same trends in GEBV, coancestry, and ultimate haplotype scores as truncation selection, but on a different (lower) scale. It loses useful diversity much faster than the

corresponding truncation selection condition (Figure 8c), even while maintaining slightly lower coancestry (Figure 8b), so its increase in GEBV before its plateau is approximately half of what truncation selection achieves (Figure 8a). Both the OCS condition matched to OHS selection's initial rate of coancestry increase and the one matched to OHS selection's initial rate of genetic gain show the same steady but slowing increase in mean GEBV over the generations as the OHS conditions, but the OCS conditions have a lower rate of increase, and after 100 generations have not caught up to truncation selection's plateaued score, whereas OHS selection has surpassed it (Figure 8a). The coancestry trends from OCS flatten out more than the OHS condition's trend (Figure 8b). The explanation for their lower improvements in GEBV compared to OHS can be seen in Figure 8c), where it is shown that the OCS conditions lose useful diversity slower than truncation selection but much faster than the GA. By the end of the 100 generations, they have lower upper selection limits than truncation selection, so do not have the potential to surpass the truncation selection GEBV plateau, even if more generations were simulated. Later sections of this paper will discuss ideas on why OCS performs poorly in this situation.

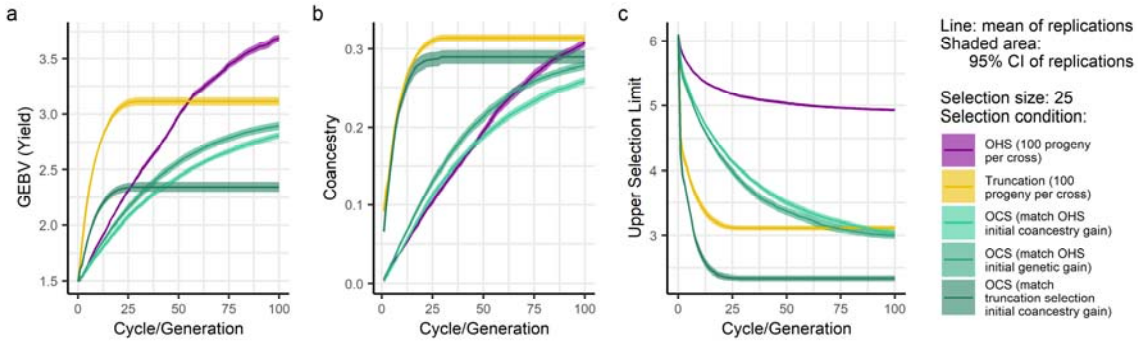

Fig. 8: Trends in simulation results, comparing optimal haplotype stacking (OHS), truncation selection, and optimal cross selection (OCS). (a) Trends in each generation's mean GEBV. (b) Trends in each generation's genetic diversity (Nei (1973)'s expected heterozygosity measure). (c) Trends in each generation's upper selection limit. Results are from simulation with a recurrent selection size of 25. Represented in the plots are the following conditions: truncation selection with 100 progeny per cross; selection by OHS with 100 progeny per cross; OCS with target angle chosen to match the initial increase in coancestry of the truncation condition; OCS with target angle chosen to match the initial increase in coancestry of the OHS condition; and OCS with target angle chosen to match the initial gain in GEBV score of the OHS condition.

Targeting one haplotype versus targeting two haplotypes per chromosome segment for genetic gain and maintaining genetic diversity

The value and trend of the upper selection limit (Figure 9c) was similar for OPV and OHS. As expected, the genetic diversity of the population under two-haplotype (OHS) stack selection is higher than it is under one-haplotype (OPV) stack selection (Figure 9b). Interestingly, the long term genetic gain under two-haplotype (OHS) stack selection was higher than the gain under the one-haplotype (OPV) selection condition, likely because OHS retains more genetic diversity.

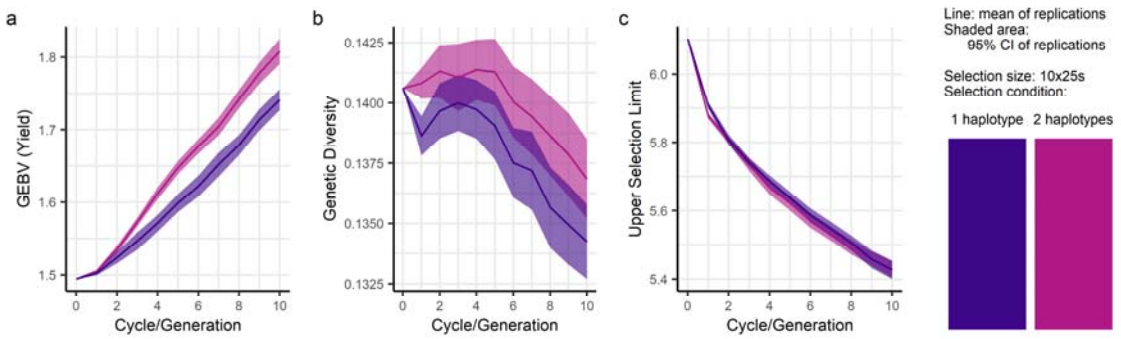

Fig. 9: Trends in 10 generations of simulation results of optimal haplotype stacking (OHS) compared to Goiffon et al. (2017)'s optimal population value (OPV), which uses a near-identical fitness function but selects on the best single haplotype at each chromosome segment where OHS selects on the best two haplotypes (from two different candidates) at each chromosome segment. (a) Trends in each generation's mean GEBV. (b) Trends in each generation's genetic diversity (Nei (1973)'s expected heterozygosity measure). (c) Trends in each generation's upper selection limit. Results are from simulation with a recurrent selection size of 25, and 10 progeny per cross.

Discussion

Our results demonstrate that OHS as a recurrent selection method maintains long term improvements and can provide greater long-term gains than either recurrent truncation selection or OCS, at least for the simulation conditions tested here. OHS shows the greatest gains over truncation selection when small populations of parents are selected and mated each year with large family sizes. For large population sizes, truncation selection was often superior to OHS in terms of genetic gain.

In our results, an appropriate mate selection strategy (OCS or OHS) can maintain genetic variation in crop breeding programs for many generations, even if the starting point is a narrow pool of elite germplasm. In the situations tested, OHS was able to achieve greater genetic gains than OCS, for OCS constrained to achieve the same coancestry as OHS (e.g. Figure 8a). The measure of coancestry incorporates genetic information, but only represents a single global score capturing the genetic difference of a pair of individuals, without accounting for the individual contribution of different molecular variants to the quantitative trait of interest. Meanwhile, OHS considers which chromosome segments differ between individuals and what their individual effect on the trait of interest is, thereby identifying useful diversity with regards to increasing the genetic merit of the target trait (Figure 8c).

The performance of OCS in this study's simulations is poorer than expected. While, for the right target angle, recurrent OCS could match the initial rate of genetic gain of truncation selection, it exhausted its diversity at a much lower GEBV plateau (Figure 8a). In the simulations of Labroo & Rutkoski (2022), recurrent OCS selection was on par with recurrent truncation selection for 50 generations for low heritabilities,and lagged behind recurrent truncation selection for high heritabilities. Here we have simulated a high heritability, by using GEBV as effectively true breeding values. Like the simulations in this paper, Labroo & Rutkoski (2022) produced a single progeny from each pair in the OCS crossing plan. Simulation of recurrent truncation selection and OCS on closed elite populations in Sanchez et al. (2023), however, showed OCS maintaining rates of genetic gain on par with truncation selection, and continuing to improve after generation 20 even as the rates of gain of truncation selection fell off and reached a plateau. Even after 60 generations, the OCS populations had not exhausted their genetic diversity. Allier et al. (2019)'s simulations also showed OCS and derivatives outperforming selection on true breeding values over 60 generations. This study's simulation show a similar number of generations for truncation

selection conditions to reach their plateau, but the OCS conditions in this study exhaust their diversity much faster than they do in Allier et al. (2019). Allier et al. (2019) produced doubled haploids from full-sib families from each suggested mating, and Sanchez et al. (2023) produced 80 doubled haploids each generation from each of the 20 pairs in the mating plan, rather than producing only one progeny for each suggested pair in the mating plan, as was done in this study. This study's results potentially do not show OCS at its best-performing.

Usually, in a breeding program, underperforming candidates would be culled before selection, whether by breeders or by failure to survive. This kind of culling was not present in this study's simulation design. Goiffon et al. (2017) found that OPV often outperformed OHS in 10 generations of selection, but our study's simulation results had OHS outperforming OPV over the same measure and time-frame (Figure 9a). Goiffon et al. (2017)'s simulations included an extra parameter representing the percentage of the population (the percentage with the lowest GEBV) that would be culled before selection. They found that 70% culling rate for OPV and 80% for OHS produced the highest genetic gains after 10 generations. The difference may be due to different founder populations in the two studies, or it may be that the pre-selection culling of low GEBV candidates conducted in Goiffon et al. (2017) was important to optimise gain under OPV.

It cannot be said whether OPV would be more appropriate than OHS for a breeding program aiming to produce inbreds or doubled haploids, because this study's simulated breeding programs did not test that kind of design. However, the fact OHS had higher population improvement on average than OPV in this study's simulations suggests that OHS would be at minimum worth investigating alongside OPV. OHS maintained more diversity than OPV (Figure 9b), so it may be preferred for breeding goals where diversity is beneficial (eg long-term selection programs, genetic resource exploration). A hybrid strategy would also be worth investigating, where several generations of OHS were followed by generations of OPV to produce inbred lines.

Gorjanc et al. (2016) discuss the problem of 'masking' in pre-breeding programs, where desirable chromosome segments (e.g. chromosome segments with favourable effects on the target trait) are contained in the genome of candidates with an overall low GEBV, and so are hidden to whole-genome-score selection measures. The OHS selection method addresses this masking problem by targeting segments as selection units directly, and by selecting at the set-of-candidates level rather than at the individual-candidate level. Mediocre-GEBV candidates can be selected because they hold good haplotype segments not present in other selection candidates with higher GEBV. Even if a candidate genotype contains only one good segment, OHS still has the chance to select it, if that good segment is very high-performing or good haplotypes for that chromosome segment are rare in the breeding pool. Allier et al. (2020) suggested that the task of initially choosing a set of donors most compatible to a set of elite lines for a pre-breeding program could be a good fit for an extension to set-of-candidates selection of the direct stacking method PCV (Han et al. 2017). OHS is an indirect stacking method (which does not consider the probability of stacking its selected segments the way that PCV does), but it is used to select sets of candidates that are compatible in a segment-stacking sense. Therefore, OHS matches the approach that Allier et al. (2020) suggested might suit this problem.

A distinguishing feature of OHS in our results was how much genetic diversity was maintained compared to other selection methods. Truncation selection conditions saw a sharp initial decline in diversity, while the trend of OHS conditions' genetic diversity measures was a much shallower, near-linear decrease. This may be because OHS's set-of-candidates and segment-based selection approach can choose to include some mediocre

candidate genotypes because they have one or few good segments, candidates which would not be selected by truncation selection. The other, mediocre, segments of such candidates will likely be different to the higher-value segments of other selected candidates, which will add to measurements of genetic diversity.

OHS has been shown to, even in one generation of application, identify useful recombinants from the offspring pool and ensure that good segments that would otherwise have been lost under truncation selection remain in the population (eg. The hybrid selection condition in Figure 5). It could therefore be well-suited to early generations after the introduction of new diversity to a breeding program, to facilitate combining new and old alleles, especially when new alleles are carried by individuals with comparatively low overall GEBV. It may also perform well in tasks involving breeding for multiple traits (e.g. stacking disease resistances), breeding for a new complex trait, or targeted exploration of gene bank germplasm resources.

The relative performance of OHS improved in both the long and short term with higher numbers of progeny per cross. This is not because OHS selects many similar genotypes. Full- or half-siblings, no matter how high their GEBVs, are likely to get that GEBV from many of the same high-scoring segments, and therefore OHS's fitness function will not select them together. Higher numbers of progeny per cross do however mean greater chances of desired recombination between favourable alleles occurring, so more chance of there being a superior recombination to select from the pool.

Improvements in performance from more progeny per cross in OHS are diminishing, with the first tenfold increase from 1 to 10 progeny per cross improving performance more than the second tenfold increase from 10 to 100 progeny per cross (Figure 5a). A balance between OHS performance improvement and the increased genotyping costs and diminishing returns of more offspring per cross must exist. But overall, its performance suggests that OHS would be most powerful in species where large numbers of progeny of the same cross can be generated and selected between, and affordable genotyping platforms exist.

Larger selection sizes mean it is possible to retain a wider variety of genotypes with more diverse and useful alleles in the population. Both truncation selection and OCS perform better with larger selection sizes, because they cannot over-cull as strongly and so will maintain more diversity to be turned into genetic gain. However, the genetic gain achieved by OHS over the time horizons tested was relatively higher for smaller populations, for example 10 or 25 parents each generation (Figure 7). This is likely because recurrent OHS has an indefinite future horizon. In a bottleneck scenario, its ultimate genotype's segments must come from a limited number of candidates, which makes it more likely that recombination events will occur that bring those good segments together, and so the observed rate of gain will be higher. Given instead a large selection size, OHS can pick exceptional segments isolated in otherwise low-value genotypes. It then requires more rare recombination events to stack the segments and see genetic gain. This behaviour is visible in Figure 7: the gap between the mean generation-100 GEBV and the score of the upper selection limit is larger for larger population sizes.

The effect of increasing the number of chromosome segments for OHS is very similar to the effect of increasing the population size. Having more chromosome segments allows OHS to be more fine-grained and ambitious in the stacking events required to create the ultimate genotype, and therefore results in very slow rates of improvement of the population, whereas having fewer larger segments allows for faster rates of gain. However, the genetic merit of the ultimate genotype is reduced for large block sizes if there are multiple QTL per segment. OPV (Goiffon et al. 2017), the method very similar to OHS (in that it also selects a set of

candidates, though it maximises the upper selection limit of the doubled haploid, choosing one rather than two haplotypes for each chromosome segment), shows the same behaviour: more segments means a longer-term selection horizon (Goiffon et al. 2017). The effects overlapping segments would have is not known.

In this, block collection methods like OHS (and OPV) are quite arbitrary as selection methods. Selection pressure is roughly controlled by the number of candidates it is set to select, and the number of segments into which the genome is divided. A method for optimising these parameters for particular time frames or desired rates of gain has not yet been developed but could potentially greatly enhance the efficiency of these methods.

Further development of the class of haplotype block collection selection strategies like OPV and OHS would not be limited to the comparison and design of different fitness function targets. Currently, marker effects are estimated with the assumption they will be used to calculate a candidate's whole-genome GEBV. Potentially epistasis and G x E interactions could be modelled differently, at the block-level, so as to be beneficially incorporated into a block-based selection metric.

Other selection methods that do not select a whole pool of candidates but instead specific mating pairs (eg PCV) or specific candidates (eg OHV) lose the ability to choose candidates based on what is already represented in the pool, and so likely lose the specific advantages in scenarios around integration of donor and elite lines. In Goiffon et al. (2017) OHV was very frequently outperformed by OHS and OPV, over 10 generations of selection. OCS and similar methods (GOCS, Genomic Mating, UCPC) do have the ability to take sets of candidates into account, and also have the probable advantage over OHS that they are easy to target for specific rate of genetic improvement or coancestry increase balance goals. It remains to be seen if these methods could, in a simulation design better suited to them, maintain useful diversity to comparable levels to OHS. Even if they are more appropriate, economically, for practical breeding programs, they may not have the single-generation application performance boost of OHS observed in Figure 4 and Figure 5a.

The goal of OHS is to identify a subset of candidates that contain, between them, chromosome segments to construct an optimal "ultimate genotype". There could be several subsets of candidates that construct similar-scoring ultimate genotypes but share no or very few members, if different candidate genotypes carry different combinations of good chromosome segments. This would create a multi-peak landscape for a GA to search. Furthermore, each candidate contributes a different amount to the ultimate genotype score of different subsets based on the segments provided by other genotypes in the subset, which results in noise in the fitness landscape. Among evolutionary algorithms, GAs perform particularly well on complex or noisy fitness landscapes (Mitchell 1998).

As an evolutionary algorithm, a GA iteratively improves its optimal subset guesses to get a close to optimal subset in a reasonable time frame. While the globally optimal subset may be a slightly better set of chromosome segments, distributed through the selection, the genetic improvement of the breeding program, especially in the short- and medium- term, depends more on how many useful chance recombination events that bring good segments together into the same genotype can occur in the allotted time and cost space. A GA therefore seems a good choice of search tool for OHS's fitness function.

The limitations of our simulations are worth pointing out. SNP effects were assumed to be true effects - additive effects that held constant over time. That is, it was assumed markers perfectly tracked the QTL over 100 generations of simulation; the controlling QTL had only additive effects; the positions and order on the chromosome of all markers were perfectly

known; and for the purposes of calculating local GEBVs for OHS, phase could be perfectly determined. QTL effects could have been simulated, however this has the issue that the number of QTL and the distribution of their effects would have to be assumed. In a real-world situation, the estimates of marker effects and estimates of phase could be iteratively updated with phenotype data collected in each new season's trials.

Further, no mutation was simulated, even though beyond 20 generations of artificial selection, the contribution of mutation to genetic variance becomes quite visible (Dudley and Lambert 2004; Barton and Keightley 2002; Keightley 2004). If marker effects that capture the new mutation's value or updated SNP arrays were available, truncation selection and OCS conditions would be able to improve past their plateaus. However, simulating mutation is quite challenging because the distribution of mutation effects is poorly defined. Future studies will investigate the effect of mutation and erosion of linkage disequilibrium between markers and QTL over time.

One of the assumptions built into the simulation tool was that crossovers occur with uniform frequency along the length of each chromosome. This means the effects of linkage disequilibrium (LD) on chances of recombination at the desired points were not investigated. Dividing chromosome segments by known LD blocks rather than by chromosome length would therefore not have provided any improvements in stacking performance in simulation. The OHS fitness function also assumes the chance of recombination between segments is uniform across all segments: though recombination in simulation occurred within segments as well as between segments, the performance of OHS when its uniformity-of-recombination assumption was properly challenged was not tested. Building into the GA's objective function some measure of this probability, for example from LD block structures, could improve the validity of its suggestions for real organisms.

Since all selection in the simulations was on true breeding values, the effect of uncertainties and performance with lower-heritability traits is not known. In reality, the predicted effects of rare haplotype blocks are unlikely to be accurate, so the likelihood of identifying a rare and beneficial haplotype block and preserving it in the population until it recombines with other good segments, would be low.

Our time horizon in simulation was 100 generations. Although this seems very long, it is worth pointing out with the advent and widespread adoption of speed breeding (Watson et al. 2018), with which up to seven generations of wheat can be turned over each year, 100 generations equates to just 15 years. In a two-part breeding program design, this time frame is both feasible and desirable in the population improvement component (which ideally does not have to slowed down by field trials or selfing generations, unlike the variety development component). Consideration of methods that, unlike the recurrent high-accuracy truncation selection shown in this study's simulations, do not risk exhausting their populations' standing diversity in relatively low numbers of years, will become important.

The simulation design was not an attempt to present an equivalent to real breeding programs. Economic equivalences between different conditions and parameter choices were also not considered. The design served to show that OHS could maintain diversity and rates of gain over very long time frames and without introgression of new diversity, how the long-term selection limit plateau was affected by one generation of OHS, and which population sizes and fullsib family sizes suited different selection methods. This study's results advise that OHS could be investigated further for practical scenarios characterised by extracting useful selections from diverse candidate pools with large breeding value differences, or recurrent selection on small candidate pools. OHS as presented in this paper might suit a scenario like the population improvement component of a two-part breeding program

(Gaynor et al. 2017), which would like to extract value from the breeding population for later development into (inbred or hybrid) products.

A practical result from this study's simulations is that a single generation of selection with OHS produces a noticeable improvement not only in the long-term performance of recurrent truncation selection, but also will surpass the genetic gain of a population that did not have that single generation of OHS within 10 generations. This benefit from a single generation of application might be of interest to breeders.

Conclusion

OHS is a selection target which provides steady rates of genetic improvement even over long time horizons, by maintaining useful segments and capturing them as they gradually stack. Even one generation of OHS selection in a recurrent truncation selection program has been shown to identify useful recombinants and ensure that good segments that would otherwise have been lost remain in the breeding population, achieving returns of higher genetic gain in simulation within 10 generations compared to pure truncation selection. OHS could therefore be suited to early generations of an elite breeding program after the introduction of new diversity (to take best advantage of the new diversity and beneficial recombinants), in resistance breeding (for stacking multiple resistances together), or in pre-breeding programs that explore genetic resources from gene banks (as it can take into account useful diversity at block level), though further investigation would be required to find how to best apply the method in these practical scenarios. Simulation results have shown that the block collection approach does an excellent job of maintaining the diversity of a closed breeding population, and in the long-term (100 generations) achieves genetic gain comparable to or in some cases exceeding truncation selection and OCS. It produces its highest rates of gain in small breeding populations (10 parents or so) with large family sizes. In this study's simulation design, OHS outperformed the similar method OPV (that only tries to capture one haplotype per chromosome segment, rather than two). This result differs from the results observed in previous study of OPV, but more investigation would be necessary to determine what characteristics of the population and of the breeding program suit OHS better than OPV. This study therefore shows that this evolutionary computing-based block collection approach to selection is a promising avenue for further investigation.

List of References

Akdemir D & Sánchez JI (2016) Efficient Breeding by Genomic Mating. *Front. Genet*. 7:210. https://doi.org/10.3389/fgene.2016.00210

Allier A, Lehermeier C, Charcosset A, Moreau L & Teyssèdre S (2019) Improving Short- and Long-Term Genetic Gain by Accounting for Within-Family Variance in Optimal Cross-Selection. *Front. Genet.* 10:1006. https://doi.org/10.3389/fgene.2019.01006

Allier A, Teyssèdre S, Lehermeier C, Charcosset A, & Moreau L (2020). Genomic prediction with a maize collaborative panel: identification of genetic resources to enrich elite breeding programs. *Theor Appl Genet* **133**, 201–215. https://doi.org/10.1007/s00122-019-03451-9

Barton NH & Keightley PD (2002). Understanding quantitative genetic variation. *Nat. Rev. Genet.* 3: 11–21. https://doi.org/10.1038/nrg700

Bernardo R (2021). Predictive breeding in maize during the last 90 years. *Crop Science*, 61:2872– 2881. https://doi.org/10.1002/csc2.20529

Cole JB & VanRaden PM (2011). Use of haplotypes to estimate Mendelian sampling effects and selection limits. *Journal of Animal Breeding and Genetics, 128*: 446-455. https://doi.org/10.1111/j.1439-0388.2011.00922.x

Cooper M, Messina CD, Podlich D, Totir LR, Baumgarten A, Hausmann NJ, Wright D & Graham G (2014). Predicting the future of plant breeding. Complementing empirical evaluation with genetic prediction. *Crop Pasture Sci 65*(4):311. https://doi.org/10.1071/cp14007

Cowling WA, Castro-Urrea FA, Stefanova KT, Li L, Banks RG, Saradadevi R, Sass O, Kinghorn BP, Siddique KHM (2023). Optimal Contribution Selection Improves the Rate of Genetic Gain in Grain Yield and Yield Stability in Spring Canola in Australia and Canada. *Plants* 12(2):383. https://doi.org/10.3390/plants12020383

Cowling WA, Li L, Siddique KHM, Henryon M, Berg P, Banks RG, Kinghorn BP (2017). Evolving gene banks: improving diverse populations of crop and exotic germplasm with optimal contribution selection, *Journal of Experimental Botany*, (68)8:1927–1939, https://doi.org/10.1093/jxb/erw406

Cowling WA, Li L, Siddique KHM, Banks RG & Kinghorn BP (2019). Modeling crop breeding for global food security during climate change. *Food and Energy Security, 8*(2), e00157. https://doi.org/10.1002/fes3.157

Crossa J, Pérez P, Hickey J, Burgueño J, Ornella L, Cerón-Rojas J, Zhang X, Dreisigacker S, Babu R, Li Y, Bonnett D, & Mathews K (2014). Genomic prediction in CIMMYT maize and wheat breeding programs. *Heredity (Edinb). 112*(1):48-60. https://doi/10.1038/hdy.2013.16

Daetwyler HD, Hayden MJ, Spangenberg GC, & Hayes BJ (2015). Selection on Optimal Haploid Value Increases Genetic Gain and Preserves More Genetic Diversity Relative to Genomic Selection, *Genetics*, *200*(4):1341–1348, https://doi.org/10.1534/genetics.115.178038

De Beukelaer H, Badke Y, Fack V, & De Meyer G (2017). Moving Beyond Managing Realized Genomic Relationship in Long-Term Genomic Selection, *Genetics, 206*(2), 1127–1138, https://doi.org/10.1534/genetics.116.194449

Dekkers JCM & Hospital F (2002). The use of molecular genetics in the improvement of agricultural populations. *Nat. Rev. Genet., 3*:22-32, https://doi.org/10.1038/nrg701

Dudley JW & Lambert RJ (2004). 100 Generations of selection for oil and protein in corn. *Plant Breed. Rev. 24*(1):79-110. https://doi.org/10.1002/9780470650240.ch5

García-Ruiz A, Cole JB, VanRaden PM, Wiggans GR, Ruiz-López FJ, & Tassell CPV (2016). Changes in genetic selection differentials and generation intervals in US Holstein dairy cattle as a result of genomic selection. *Proceedings of the National Academy of Sciences*, 113(28):E3995-E4004. https://doi.org/10.1073/pnas.1519061113

Guinan FL, Wiggans GR, Norman HD, Dürr JW, Cole JB, Van Tassell CP, Misztal I, & Lourenco D (2023). Changes in genetic trends in US dairy cattle since the implementation of genomic selection. *Journal of Dairy Science, 106*(2):1110-1129, https://doi.org/10.3168/jds.2022-22205

Goiffon M, Kusmec A, Wang L, Hu G, & Schnable PS (2017). Improving Response in Genomic Selection with a Population-Based Selection Strategy: Optimal Population Value Selection, *Genetics*, *206*(3):1675–1682. https://doi.org/10.1534/genetics.116.197103

Gaynor RC, Gorjanc G, Bentley AR, Ober ES, Howell P, Jackson R, Mackay IJ, & Hickey JM (2017). A Two-Part Strategy for Using Genomic Selection to Develop Inbred Lines. *Crop Science*, *57*(5), 2372-2386. https://doi.org/https://doi.org/10.2135/cropsci2016.09.0742

Gorjanc G, Gaynor RC, & Hickey JM (2018). Optimal cross selection for long-term genetic gain in two-part programs with rapid recurrent genomic selection. *Theoretical and Applied Genetics, 131*(9):1953-1966. https://doi.org/10.1007/s00122-018-3125-3

Gorjanc G & Hickey JM (2018). AlphaMate: a program for optimising selection, maintenance of diversity, and mate allocation in breeding programs. *Bioinformatics*, 34(19):3408-3411, https://doi.org/10.1093/bioinformatics/bty375

Gorjanc G, Jenko J, Hearne SJ, & Hickey JM (2016). Initiating maize pre-breeding programs using genomic selection to harness polygenic variation from landrace populations. *BMC Genomics, 17*(1):30. https://doi.org/10.1186/s12864-015-2345-z

Han Y, Cameron JN, Wang L, & Beavis WD (2017). The Predicted Cross Value for Genetic Introgression of Multiple Alleles, *Genetics*, *205*(4), 1409–1423, https://doi.org/10.1534/genetics.116.197095

Hickey J, Chiurugwi T, Mackay I, Powell W & Implementing Genomic Selection in CGIAR Breeding Programs Workshop Participants (2017). Genomic prediction unifies animal and plant breeding programs to form platforms for biological discovery. *Nat Genet* **49**:1297–1303. https://doi.org/10.1038/ng.3920

Hospital F (2003). Marker-assisted breeding. H.J. Newbury (Ed.), Plant Molecular Breeding, Blackwell Scientific Publishers, London, pp. 30-56.

Isik F, Holland J, & Maltecca C (2017). Genomic Relationships and GBLUP. *Genetic Data Analysis for Plant and Animal Breeding*. Springer, Cham. https://doi.org/10.1007/978-3-319-55177-7_11

Keightley PD (2004). Mutational variation and long-term selection response. *Plant Breed. Rev. 24*(1):227-247. https://doi.org/10.1002/9780470650240.ch10

Kemper KE, Bowman PJ, Pryce JE, Hayes BJ, & Goddard ME (2012). Long-term selection strategies for complex traits using high-density genetic markers. *Journal of Dairy Science, 95*(8):4646-4656. https://doi.org/10.3168/jds.2011-5289

Kerr RJ, McRae TA, Dutkowski GW, & Tier B (2015). Managing the rate of increase in average co-ancestry in a rolling front tree breeding strategy. *Journal of Animal Breeding and Genetics, 132*(2):109-120. https://doi.org/10.1111/jbg.12157

Kinghorn BP (2011). Automated Mate Selection Analyses. *Proc. Assoc. Advmt. Anim. Breed. Genet., 19*:15-22.

Labroo MR & Rutkoski JE (2022). New cycle, same old mistakes? Overlapping vs. discrete generations in long-term recurrent selection. *BMC Genomics, 23*, 736. https://doi.org/10.1186/s12864-022-08929-3

Lin Z, Cogan NOI, Pembleton LW, Spangenberg GC, Forster JW, Hayes BJ, & Daetwyler H D (2016). Genetic Gain and Inbreeding from Genomic Selection in a Simulated Commercial Breeding Program for Perennial Ryegrass. *The Plant Genome, 9*(1), https://doi.org/10.3835/plantgenome2015.06.0046

Meuwissen TH (1997). Maximizing the response of selection with a predefined rate of inbreeding. *J Anim Sci., 75*(4):934-40. https://doi.org/10.2527/1997.754934x

Meuwissen TH, Hayes BJ, Goddard ME (2001). Prediction of total genetic value using genome-wide dense marker maps. *Genetics 157*(4):1819–1829. https://doi.org/10.1093/genetics/157.4.1819

Mitchell M (1998). *An introduction to genetic algorithms* (1st MIT Press paperback ed.). MIT Press.

Nei M (1973). Analysis of gene diversity in subdivided populations*. Proc. Natl. Acad. Sci. USA 70:3321–3323.* https://doi.org/10.1073/pnas.70.12.3321

Saradadevi R, Mukankusi C, Li L, Amongi W, Mbiu JP, Raatz B, Ariza D, Beebe S, Varshney RK, Huttner E, Kinghorn B, Banks R, Rubyogo JC, Siddique KHM, & Cowling WA (2021). Multivariate genomic analysis and optimal contributions selection predicts high genetic gains in cooking time, iron, zinc, and grain yield in common beans in East Africa. *Plant Genome*, 14:e20156. https://doi.org/10.1002/tpg2.20156

Sanchez D, Sadoun SB, Mary-Huard T, Allier A, Moreau L, & Charcosset A (2023). Improving the use of plant genetic resources to sustain breeding programs' efficiency. *Proceedings of the National Academy of Sciences*, *120*(14), e2205780119. https://doi.org/doi:10.1073/pnas.2205780119

Servin B, Martin OC, Mézard M, Hospital F (2004). Toward a theory of marker-assisted gene pyramiding. *Genetics, 168*(1):513-23. https://doi.org/10.1534/genetics.103.023358

Scott BA, Haile-Mariam M, Cocks BG, & Pryce JE (2021). How genomic selection has increased rates of genetic gain and inbreeding in the Australian national herd, genomic information nucleus, and bulls. *Journal of Dairy Science 104*(11):11832-11849, https://doi.org/10.3168/jds.2021-20326

Tiret M, Pégard M, & Sánchez L (2021). How to achieve a higher selection plateau in forest tree breeding? Fostering heterozygote × homozygote relationships in optimal contribution selection in the case study of Populus nigra. *Evolutionary Applications, 14*, 2635–2646. https://doi.org/10.1111/eva.13300

Villiers K, Dinglasan E, Hayes BJ, & Voss-Fels KP (2022). genomicSimulation: fast R functions for stochastic simulation of breeding programs, *G3 Genes|Genomes|Genetics*, 12(10), jkac216, https://doi.org/10.1093/g3journal/jkac216

Voss-Fels KP, Cooper M, Hayes BJ (2019). Accelerating crop genetic gains with genomic selection. *Theor Appl Genet. 132*(3):669-686. https://doi.org/10.1007/s00122-018-3270-8

Watson A, Ghosh S, Williams MJ, Cuddy WS, Simmonds J, Rey M-D, Asyraf Md Hatta M, Hinchliffe A, Steed A, Reynolds D, Adamski NM, Breakspear A, Korolev A, Rayner T, Dixon LE, Riaz A, Martin W, Ryan M, Edwards D, Domoney C, Uauy C, Hazard B, Wulff BBH, & Hickey LT (2018). Speed breeding is a powerful tool to accelerate crop research and breeding. *Nature Plants*, *4*(1), 23-29. https://doi.org/10.1038/s41477-017-0083-8

Wolc A, Zhao HH, Arango J, Settar P, Fulton JE, O'Sullivan NP, Preisinger R, Stricker C, Habier D, Fernando RL, Garrick DJ, Lamont SJ, and Dekkers JCM (2015). Response and inbreeding from a genomic selection experiment in layer chickens. *Genet Sel Evol* **47**:59. https://doi.org/10.1186/s12711-015-0133-5

Woolliams JA, Berg P, Dagnachew BS, & Meuwissen THE (2015). Genetic contributions and their optimization. *Journal of Animal Breeding and Genetics, 132*(2):89-99. https://doi.org/10.1111/jbg.12148

Wray NR & Goddard ME (1994). Increasing long-term response to selection. *Genetics, Selection, Evolution : GSE, 26*:431-451, https://doi.org/10.1186/1297-9686-26-5-431

Statements and Declarations

Funding

authors are grateful for funding from the Australian Research Council, Linkage Project LP170100317, “Fastack - evolutionary computing to stack desirable alleles in wheat”.

Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

Author Contributions

All authors contributed to the study conception and design. The first draft of the manuscript was written by KV and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Data Availability

The simulated datasets generated during the current study are available from the corresponding author on reasonable request.
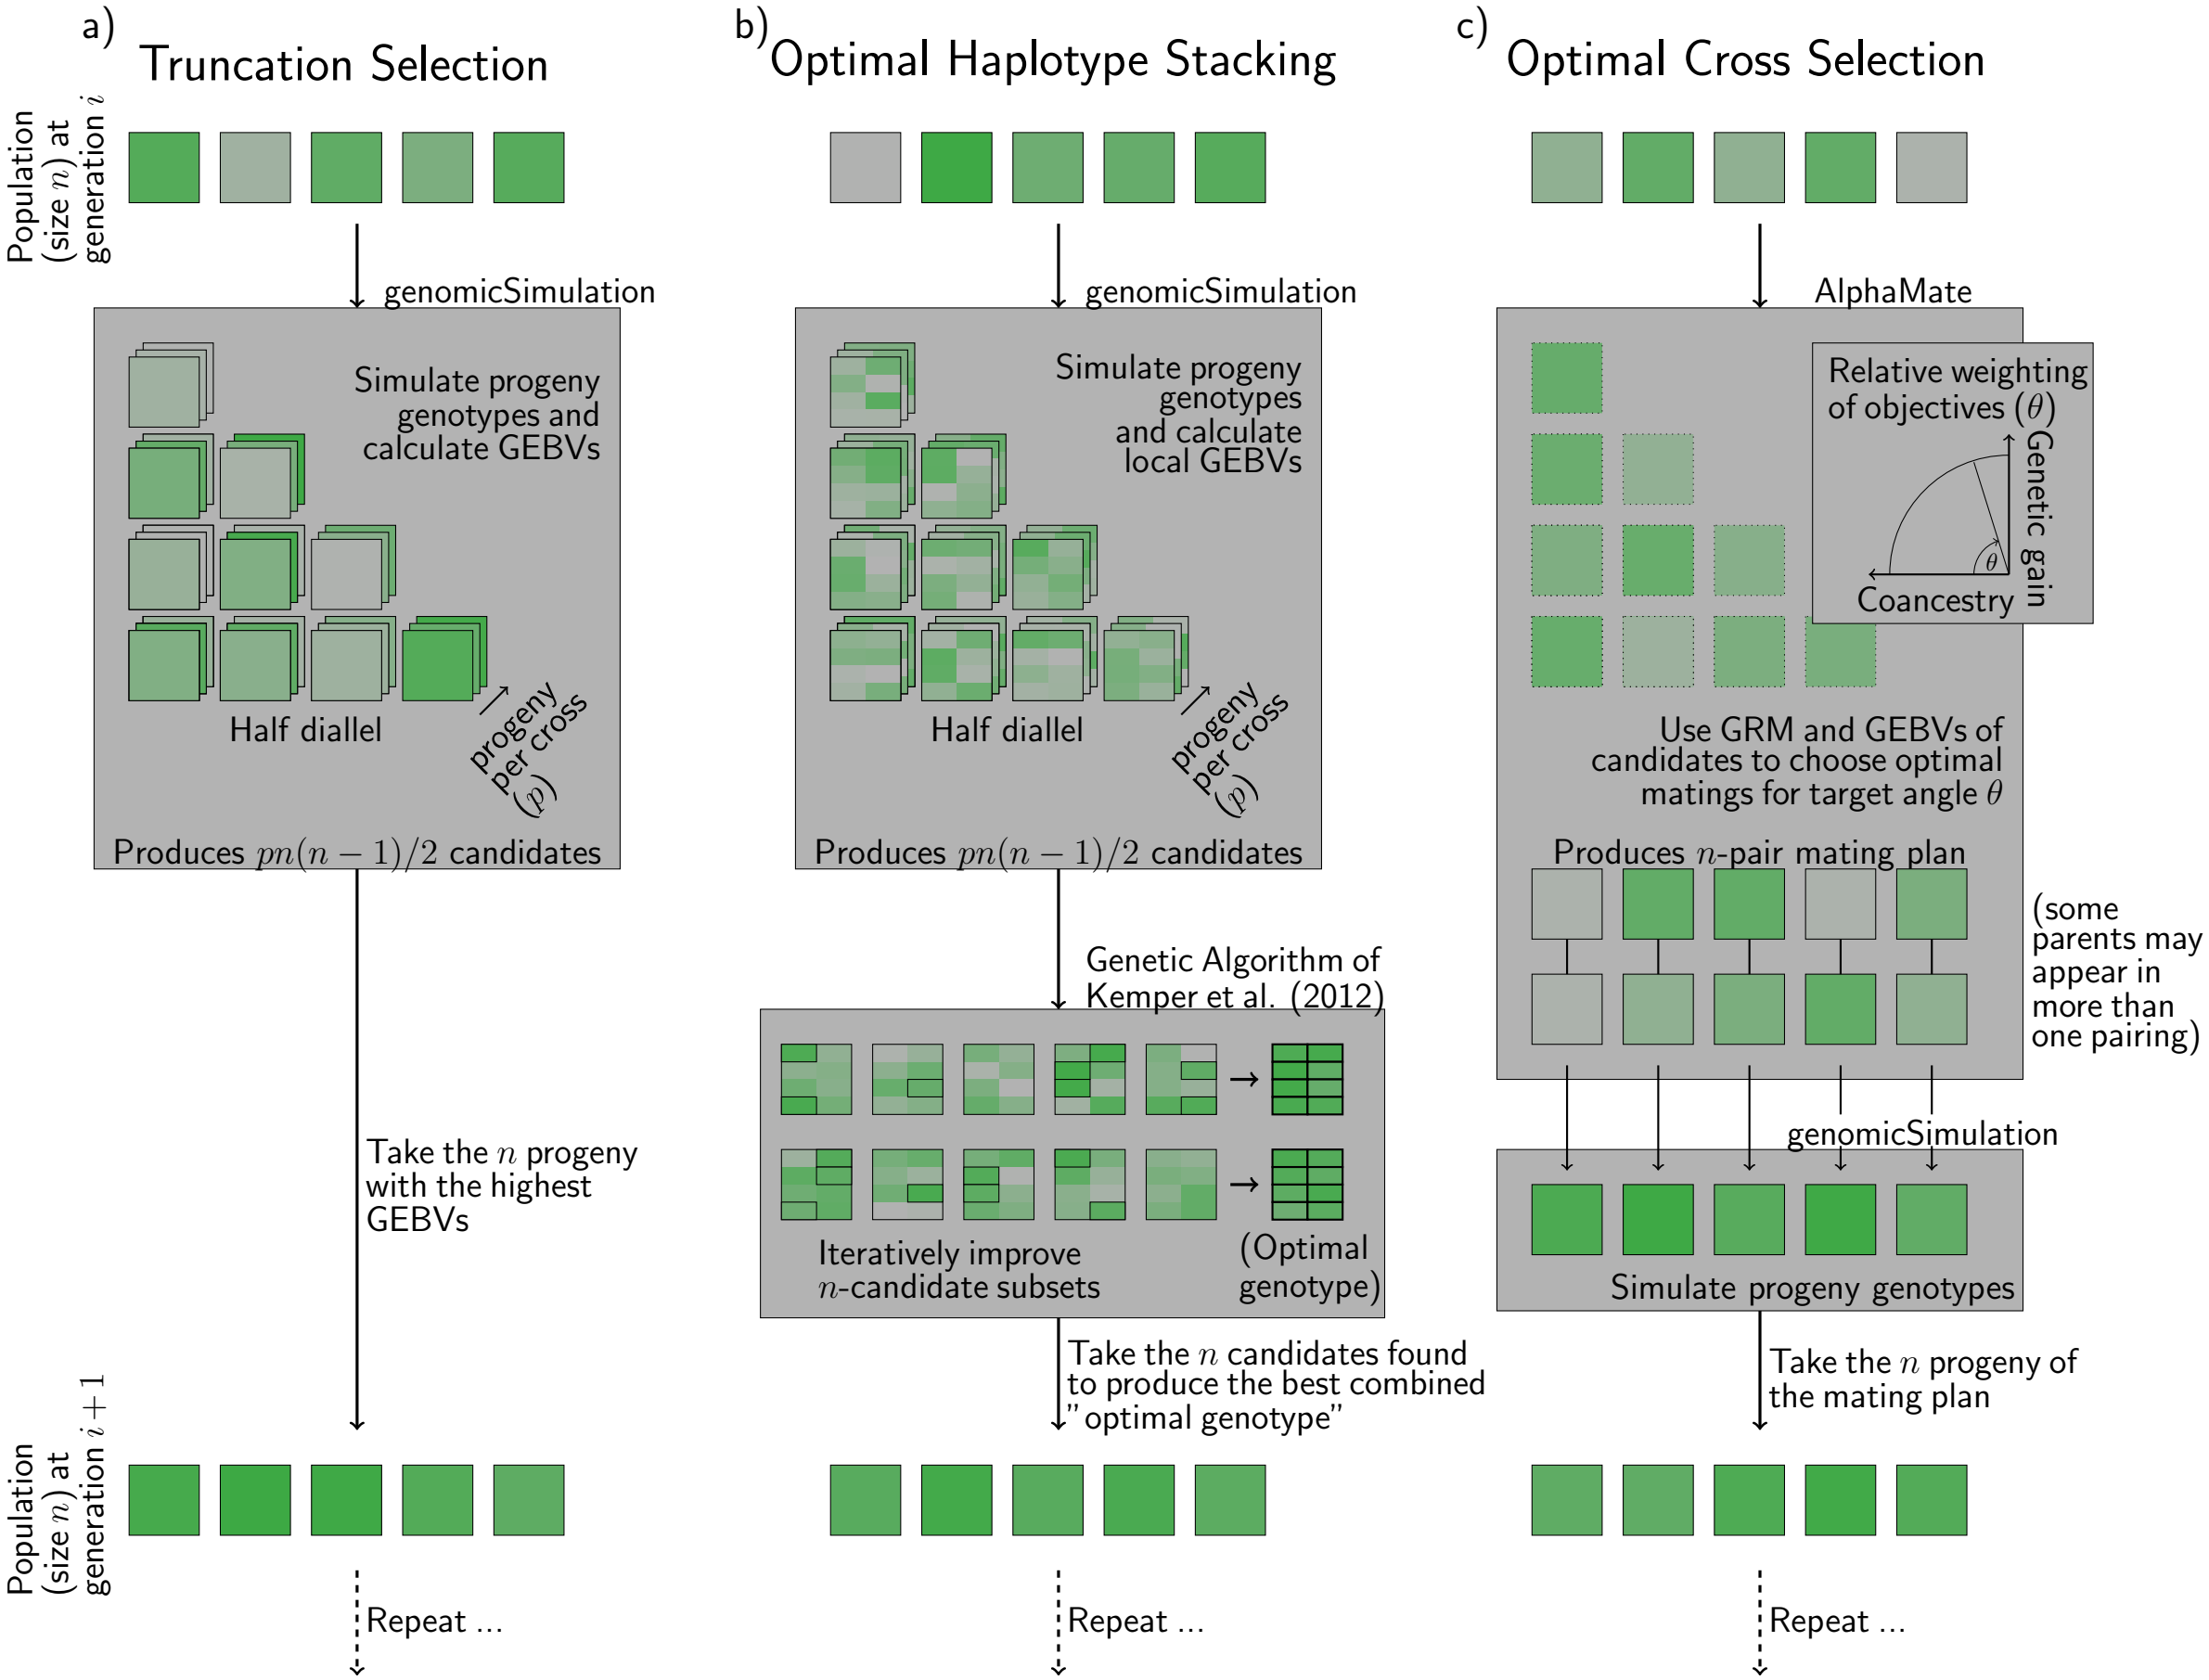
Simulation of recurrent ...