

562 **A Consistency of Bounce**

563 In this section, we prove the consistency of the Bounce algorithm. The proof is based on Papenmeier
564 et al. [48] and Eriksson and Poloczek [20].

565 **Theorem 1** (Bounce consistency). *With the following definitions*

566 *Def. 1. $(\mathbf{x}_k)_{k=1}^\infty$ is a sequence of points of decreasing function values;*

567 *Def. 2. $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}}$ is a minimizer of f in \mathcal{X} ;*

568 *and under the following assumptions:*

569 *Ass. 1. D is finite;*

570 *Ass. 2. f is observed without noise;*

571 *Ass. 3. The range of f is bounded in \mathcal{X} , i.e., $\exists C \in \mathbb{R}_{++}$ s.t. $|f(\mathbf{x})| < C \forall \mathbf{x} \in \mathcal{X}$;*

572 *Ass. 4. For at least one of the minimizers \mathbf{x}_i^* the (partial) assignment corresponding to the continu-
573 ous variables lies in a (continuous) region with positive measure;*

574 *Ass. 5. One Bounce reached the input dimensionality D , the continuous elements of the initial
575 points $\{\mathbf{x}_{cont_i}\}_{n=1}^{n_{min}}$ after each TR restart are chosen*

576 *(a) uniformly at random for continuous variables; and*

577 *(b) such that every realization of the combinatorial variables has positive probability;*

578 *then the Bounce algorithm finds a global optimum with probability 1, as the number of samples N
579 goes to ∞ .*

580 *Proof.* The range of f is bounded per Assumption 3, and Bounce only considers a function evalua-
581 tion a ‘success’ if the improvement over the current best solution exceeds a certain constant thresh-
582 old. Bounce can only have a finite number of ‘successful’ evaluations because the range of f is
583 bounded per Assumption 3. For the sake of a contradiction, we suppose that Bounce does not obtain
584 an optimal solution as its number of function evaluations $N \rightarrow \infty$. Thus, there must be a sequence
585 of failures, such that the TRs in the current target space, i.e., the current subspace, will eventually
586 reach its minimum base length. Recall that in such an event, Bounce increases the target dimension
587 by splitting up the ‘bins’, thus creating a subspace of $(b + 1)$ -times higher dimensionality. Then
588 Bounce creates a new TR that again experiences a sequence of failures that lead to another split, and
589 so on. This series of events repeats until the embedded subspace eventually equals the input space
590 and thus has dimensionality D . See lines 12 – 16 in Algorithm 1 in Sect. 3.

591 Still supposing that Bounce does not find an optimum in the input space, there must be a sequence
592 of failures such that the side length of the TR again falls below the set minimum base length, now
593 forcing a restart of Bounce. Recall that at every restart, Bounce samples a fresh set of initial points
594 uniformly at random from the input space; see line 18 in Algorithm 1. Therefore, with probability 1,
595 a random sample will eventually be drawn from any subset $\mathcal{Y} \subseteq \mathcal{X}$ with positive Lebesgue measure
596 ($\nu(\mathcal{Y}) > 0$):

$$1 - \lim_{k \rightarrow \infty} (1 - \mu(\mathcal{Y}))^k = 1, \quad (2)$$

597 where μ is the uniform probability measure of the sampling distribution that Bounce employs for
598 initial data points upon restart [91].

599 Let

$$\alpha = \inf \{t : \nu [x \in \mathcal{X} \mid f(x) < t] > 0\}$$

600 denote the essential infimum of f on \mathcal{X} with ν being the Lebesgue measure [91].

601 Following Solis and Wets [91], we define the optimality region, i.e., the set of points whose function
602 value is larger by at most ε than the essential infimum:

$$R_{\varepsilon, M} = \{x \in \mathcal{X} \mid f(x) < \alpha + \varepsilon\}$$

603 with $\varepsilon > 0$ and $M < 0$. Because of Ass. 4, at least one optimal point lies in a region of positive
604 measure that is continuous for the continuous variables. Therefore, we have that $\alpha = f(\mathbf{x}^*)$. Note
605 that this is also the case if the domain of f only consists of combinatorial variables (Ass. 5). Then,
606 $R_{\varepsilon, M} = \{x \in \mathcal{X} \mid f(x) < f(\mathbf{x}^*) + \varepsilon\}$.

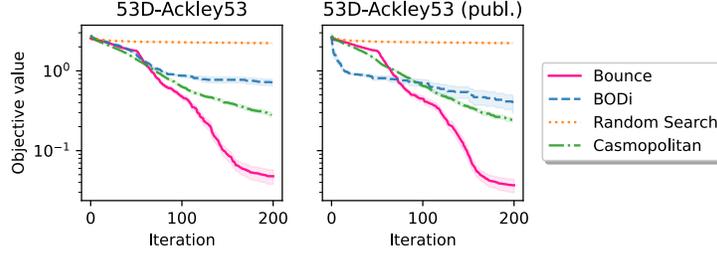


Figure 7: Bounce the other algorithms on the synthetic Ackley53 benchmark function. Bounce outperforms all other algorithms and quickly finds excellent solutions. BODi’s performance degrades upon randomization.

607 Let $(\mathbf{x}_k^*)_{k=1}^\infty$ denote the sequence of best points that Bounce discovers with \mathbf{x}_k^* being the best point
 608 up to iteration k . This sequence satisfies Def. 1 by construction. Note that $\mathbf{x}_k^* \in R_{\varepsilon, M}$ implies that
 609 $\mathbf{x}_{k'}^* \in R_{\varepsilon, M}$ for all $k' \geq k + 1$ [91] because observations are noise-free. Then,

$$\begin{aligned} \mathbb{P}[\mathbf{x}_k^* \in R_{\varepsilon, M}] &= 1 - \mathbb{P}[\mathbf{x}_k^* \in \mathcal{X} \setminus R_{\varepsilon, M}] \\ &\geq 1 - (1 - \mu(R_{\varepsilon, M}))^k, \end{aligned}$$

610 and,

$$1 \geq \lim_{k \rightarrow \infty} \mathbb{P}[\mathbf{x}_k^* \in R_{\varepsilon, M}] \geq 1 - \underbrace{\lim_{k \rightarrow \infty} (1 - \mu(R_{\varepsilon, M}))^k}_{=1, \text{ Eq. (2)}} = 1,$$

611 i.e., \mathbf{x}_k^* eventually falls into the optimality region [91]. By letting $\varepsilon \rightarrow 0$, \mathbf{x}_k^* converges to the global
 612 optimum with probability 1 as $k \rightarrow \infty$.

613 □

614 B Additional experiments

615 We compare Bounce to the other algorithms on three additional benchmark problems: Ackley53
 616 and MaxSAT60 [18]. Moreover, we run two additional studies to further investigate the performance
 617 of Bounce. First, we run Bounce on a set of continuous problems from Papenmeier et al. [48]
 618 to showcase the performance and scalability of Bounce on purely continuous problems. We then
 619 present a “low-sequency” version of Bounce to showcase how such a version can outperform its
 620 competitors on the original benchmarks by introducing a bias towards low-sequency solutions.

621 B.1 Bounce and other algorithms on additional benchmarks

622 B.1.1 The synthetic Ackley53 benchmark function

623 Ackley53 is a 53-dimensional function with 50 binary and 3 continuous variables. Wan et al. [61]
 624 discretized 50 continuous variables of the original Ackley function, requiring these variables to be
 625 either zero or one. This benchmark was designed such that the optimal value of 0.0 is at the ori-
 626 gin $\mathbf{x} = (0, \dots, 0)$. Here, we perturb the optimal assignment of combinatorial variables by flipping
 627 each binary variable with probability $1/2$. Figure 7 summarizes the performances of the algorithms.
 628 Bounce outperforms all other algorithms and proves to be robust to the location of the optimum
 629 point. Casmopolitan is a distanced runner-up. BODi initially outperforms Casmopolitan on the
 630 published benchmark version but falls behind later.

631 B.1.2 Contamination control

632 The Contamination benchmark models a supply chain with 25 stages [32]. At each stage, a binary
 633 decision is made whether to quarantine food that has not yet been contaminated. Each such inter-
 634 vention is costly, and the goal is to minimize the number of contaminated products and prevention
 635 cost [6, 45]. Figure 8 shows the performances of the algorithms.

636 Bounce, Casmopolitan, and BODi all produce solutions of comparable objective value. Bounce
 637 and Casmopolitan find better solutions than BODi initially, but after about 100 function evaluations,
 638 the solutions obtained by the three algorithms are typically on par.

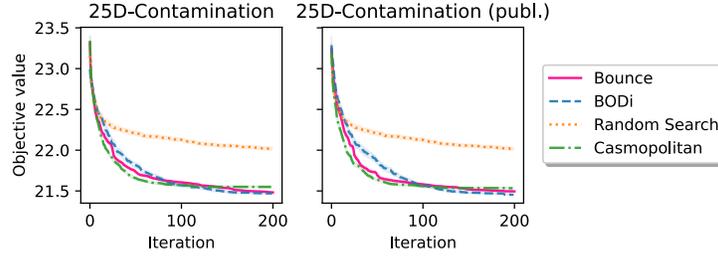


Figure 8: Bounce and the other algorithms on the 25-dimensional contamination problem. Bounce performs on par with Casmopolitan and BODi on both versions of the benchmark.

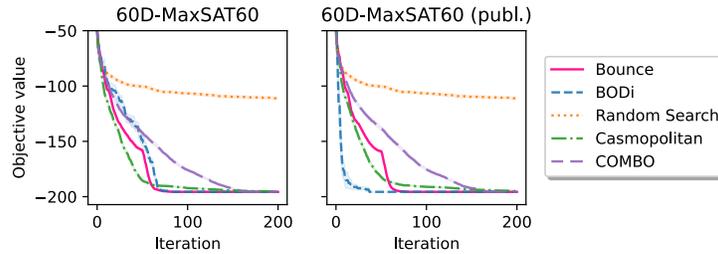


Figure 9: Bounce and other algorithms on the 60-dimensional weighted maximum satisfiability problem. Bounce is the first to find an optimal solution (left). On the published version (right), Bounce comes in second after BODi.

639 B.1.3 The MaxSAT60 benchmark

640 MaxSAT60 is a 60-dimensional, weighted instance of the Maximum Satisfiability (MaxSAT) problem.
 641 MaxSAT is a notoriously hard combinatorial problem that cannot be solved in polynomial time
 642 (unless $\mathbb{P} = \mathbb{NP}$). The goal is to find a binary assignment to the variables that satisfies clauses of maximum
 643 total weight. For every i in $[d]$, this benchmark has one clause of the form x_i with a weight of
 644 1 and 638 clauses of the form $\neg x_i \vee \neg x_j$ with a weight of 61. Following [18, 45, 61], we normalize
 645 these weights to have zero mean and unit standard deviation. This normalization causes the one-
 646 variable clauses to have a *negative weight*, i.e., the function value improves if such a clause is not
 647 satisfied, which is atypical behavior for a MaxSAT problem. Since the clauses with two variables are
 648 satisfied for $x_i = x_j = 0$ and the clauses with one variable of negative weights are never satisfied for
 649 $x_i = 0$, the normalized benchmark version has a global optimum at $\mathbf{x}^* = (0, \dots, 0)$ by construction.
 650 The problem’s difficulty is finding an assignment for variables such that *all* two-variable clauses are
 651 satisfied and *as many* one-variable clauses as possible is not captured by normalized weights.

652 Figure 9 summarizes the performances of the algorithms. The general version that attains the global
 653 optimum for a randomly selected binary assignment is shown on the left. The special case where
 654 the global optimum is set to the all-zero assignment is shown on the right.

655 We observe that Bounce requires the smallest number of samples to find an optimal assignment in
 656 general, followed by BODi and Casmopolitan. Only in the special case where the optimum is the
 657 all-zero assignment, BODi ranks first, confirming the corresponding result in Deshwal et al. [18].

658 **B.2 Parallel evaluations on continuous problems**

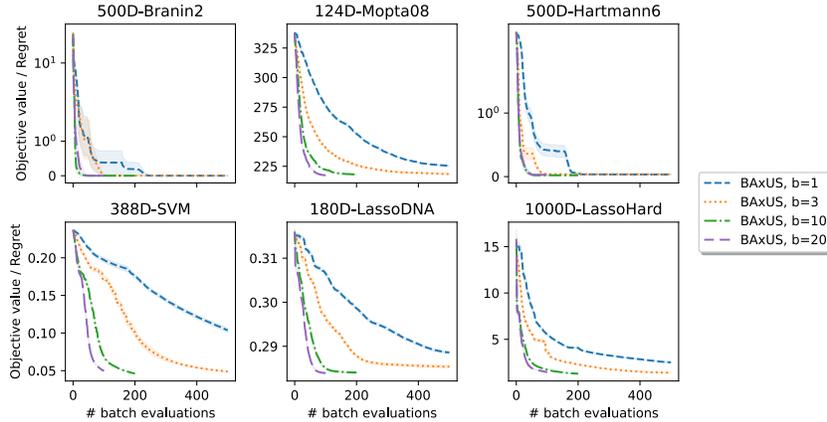


Figure 10: Bounce on continuous problems with different batch sizes.

659 To showcase the performance and scalability of Bounce, we run it on a set of continuous problems
 660 from Papenmeier et al. [48]. The 124-dimensional Mopta08 benchmark is a constrained vehicle
 661 optimization problem. We adopt the soft-constrained version from Eriksson et al. [21], Eriksson and
 662 Jankowiak [74]. The 388-dimensional soft-constrained SVM problem [74] concerns the classification
 663 performance with an SVR on the slice localization dataset. The 180-dimensional LassoDNA bench-
 664 mark [90] is a sparse regression problem on a real-world dataset, the 1000-dimensional LassoHard
 665 benchmark optimizes over a synthetic dataset. The 500-dimensional Branin2 and Hartmann6 prob-
 666 lems are versions of the 2- and 6-dimensional benchmark problems where additional dimensions
 667 with no effect on the function value were added.

668 We set the number of function evaluations to $\max(2000, 500B)$ for a batch size of B and configure
 669 Bounce such that it reaches the input dimensionality after 500 function evaluations. Figure 10 shows
 670 the simple regret for the synthetic Branin2 and Hartmann6 problems, and the best function value
 671 obtained after a given number of batch evaluations for the remaining problems: Mopta08, SVM,
 672 LassoDNA, and LassoHard.

673 We observe that Bounce always benefits from more parallel function evaluations. The difference
 674 between smaller batch sizes such as the difference between $B = 1$ and $B = 3$ or $B = 3$ and $B = 10$
 675 is more remarkable than the difference between larger batch sizes like $B = 10$ and $B = 20$. On
 676 SVM and LassoDNA, parallel function evaluations prove especially effective. Here, the optimization
 677 performance improves drastically. We conclude that a small number of parallel function evaluations
 678 already helps increase the optimization performance considerably.

679 On the synthetic Branin2 and Hartmann6 problems, Bounce quickly converges to the global opti-
 680 mum. Here, we see that a larger number of parallel function evaluations also helps in converging to
 681 a better solution.

682 **B.3 Low-sequence version of Bounce**

683 We show how we can bias Bounce towards low-sequence solutions. We do this by removing the
 684 random signs (for binary and continuous variables) and the random offsets (for categorical and
 685 ordinal variables) from the Bounce embedding. We conduct this study to show a) that Bounce is
 686 able to outperform BODi on the unmodified versions of the benchmark problems if we introduce
 687 a similar bias towards low-sequence solutions, and b) that the random signs empirically show to
 688 remove biases towards low-sequence solutions. However, we want to emphasize that the results of
 689 this section are not representative of the performance of Bounce on arbitrary real-world problems.
 690 Nevertheless, if one knows that the problem at hand has a low-sequence structure, then Bounce can
 691 be configured to exploit this structure and outperform BODi.

692 Figure 11 shows the results of the low-sequence version of Bounce on the original benchmarks from
 693 Section 4. We observe that Bounce outperforms BODi and the other algorithms on the unmodified

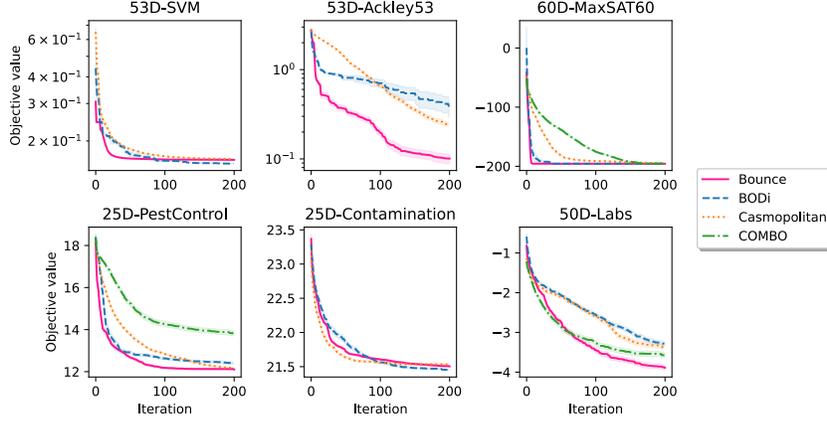


Figure 11: ‘Low-frequency’ version of Bounce on the **original** benchmarks from Section 4: with a bias towards low-frequency solutions, Bounce outperforms BODi on the original versions of the benchmark problems.

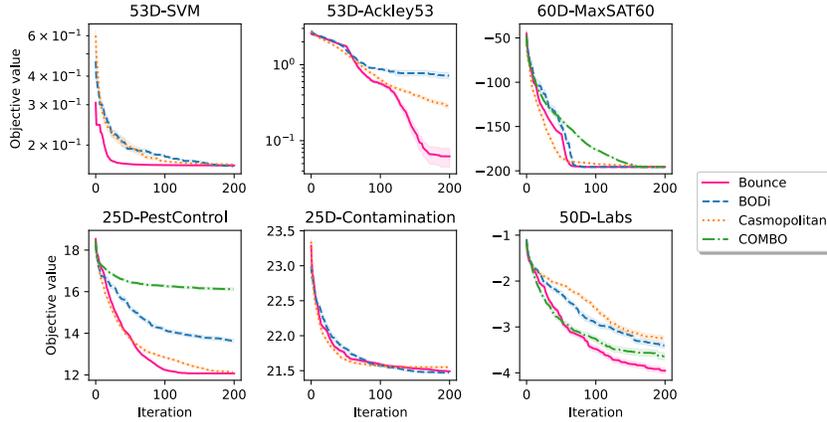


Figure 12: “Low-frequency” version of Bounce on the **modified** benchmarks from Section 4.

694 versions of the benchmark problems. This shows that Bounce can outperform BODi on the unmodi-
 695 fied version of the benchmarks if we introduce a similar bias towards low-frequency solutions.

696 Figure 12 shows the results of the low-frequency version of Bounce on the flipped benchmarks from
 697 Section 4. The low-frequency version of Bounce is robust towards randomization of the optimal
 698 point.

699 C Implementation details

700 We implement Bounce in Python using the BoTorch [5] and GPyTorch [77] libraries. For the GP
 701 model, we use as similar construction as the CoCaBo kernel [53]. We model the continuous and
 702 combinatorial variables with two separate kernels where we use automatic relevance determination
 703 (ARD) for the continuous variables and the same lengthscale for the combinatorial variables. We
 704 also simply use a Matérn kernel for the combinatorial variables. Following Ru et al. [53], we use a
 705 mixture of the sum and the product kernel:

$$k(\mathbf{x}, \mathbf{x}') = \lambda k_{\text{cmb}}(\mathbf{x}_{\text{cmb}}, \mathbf{x}'_{\text{cmb}}) k_{\text{cnt}}(\mathbf{x}_{\text{cnt}}, \mathbf{x}'_{\text{cnt}}) + (1 - \lambda)(k_{\text{cmb}}(\mathbf{x}_{\text{cmb}}, \mathbf{x}'_{\text{cmb}}) + k_{\text{cnt}}(\mathbf{x}_{\text{cnt}}, \mathbf{x}'_{\text{cnt}})),$$

706 where \mathbf{x}_{cnt} and \mathbf{x}_{cmb} are the continuous and combinatorial variables in \mathbf{x} , respectively, and λ is
 707 between 0 and 1. The kernel for the continuous variables, k_{cnt} is an ARD-Matérn kernel with
 708 $\nu = 2.5$, and the kernel for the combinatorial variables, k_{cmb} is a Matérn kernel without ARD with
 709 $\nu = 2.5$. The trade-off parameter λ is learned jointly with the other hyperparameters during the
 710 likelihood maximization.

711 We employ a $\Gamma(1.5, 0.1)$ prior on the lengthscales of both kernels and a $\Gamma(1.5, 0.5)$ prior on the
712 signal variance. We further use a $\Gamma(1.1, 0.1)$ prior on the noise variance.

713 Motivated by Wan et al. [61] and Eriksson et al. [21], we use an initial trust region baselength of 40
714 for the combinatorial variables, and 0.8 for the continuous variables. We maintain two separate TR
715 shrinkage and expansion parameters (γ_{cmb} and γ_{cnt}) for the combinatorial and continuous variables,
716 respectively such that each TR base length reaches its respective minimum of 1 and 2^{-7} after a
717 given number of function evaluations. When Bounce finds a better or worse solution, we increase
718 or decrease both TR base lengths.

719 We use the author’s implementations for COMBO¹, BODi², and Casmopolitan³. We use the same
720 settings as the authors for COMBO and BODi. For Casmopolitan, we use the same settings as the
721 authors for benchmarks reported in Wan et al. [61] and set the initial trust region base length to 40
722 otherwise.

723 Due to its high-memory footprint, we ran BODi on NVidia A100 80GB GPUs for 300 GPU/h. We
724 ran Bounce on NVidia A40 GPUs for 2,000 GPU/h. We ran the remaining methods for 20,000
725 GPU/h on one core of Intel Xeon Gold 6130 CPUs with 60GB of memory.

726 C.1 Optimization of the acquisition function

727 We use different strategies to optimize the acquisition function depending on the type of variables
728 present in a problem.

729 **Continuous problems.** For purely continuous problems, we follow a similar approach as Eriksson
730 et al. [21]. In particular, we use the lengthscales of the GP posterior to shape the TR. We use gradient
731 descent to optimize the acquisition function within the TR bounds with 10 random restarts and 512
732 raw samples. For a batch size of 1, we use analytical EI. For larger batch sizes, we use the BoTorch
733 implementation of qEI [5, 66, 88].

734 **Binary problems.** For all problems with a combinatorial search space, we use local search to
735 optimize the acquisition function. Similar to [61], we use discrete TRs around the current best
736 solution. The TR is defined as the set of all solutions that differ from the current best solution with
737 a certain Hamming distance.

738 When starting the optimization, we first create a set of $\min(5000, \max(2000, 200 \cdot d_i))$ random
739 solutions. The choice of the number of random solutions is based on Eriksson et al. [21]. For each
740 candidate, we first draw L_i indices uniformly at random from $[d_i]$ without replacement, where L_i
741 is the TR length at the i -th iteration. We then sample d_i values $\in \{0, 1\}$ and set the candidate at
742 the sampled indices to the sampled values. All other values are set to the values of the current best
743 solution. Note that this construction keeps each solution in the TR of the current best solution. We
744 add all neighbors (i.e., points with a Hamming distance of 1) of the current best solution to the
745 set of candidates. This is inspired by [18]. We find the 20 candidates with the highest acquisition
746 function value and use local search to optimize the acquisition function within the TR bounds. At
747 each local search step, we create all neighbors that do not coincide with the current best solution
748 or would violate the TR bounds. We then move the current best solution to the neighbor with the
749 highest acquisition function value. We repeat this process until the acquisition function value does
750 not increase anymore. Finally, we return the best solution found during the local search.

751 **Categorical problems.** We use the same approach as for purely binary problems, i.e., we first
752 create a set of random solutions with the same size as for purely binary problems and start the local
753 search on the 20 best initial candidates. We use one-hot encoding for categorical variables.

754 Suppose the number of categorical variables of the problem is smaller or equal to the current TR
755 length. In that case, we sample, for each candidate and each categorical variable, an index uniformly
756 at random from $[|v_i|]$ where $|v_i|$ is the number of values of the i -th categorical variable. We then set
757 the candidate at the sampled index to 1 and all other values to 0.

758 If the number of categorical variables of the problem is larger than the current TR length L_i , we first
759 sample L_i categorical variables uniformly at random from $[d_i]$ without replacement. For each initial

¹<https://github.com/QUVA-Lab/combo>, unspecified license, last access: 2023-05-04

²<https://github.com/aryandeshwal/bodi>, no license provided, last access: 2023-05-04

³<https://github.com/xingchenwan/casmo>, MIT license, last access: 2023-05-04

760 candidate and each sampled categorical variable, we sample an index uniformly at random, at which
761 we set the categorical variable to 1 and all other values to 0. The values for the variables that were
762 not sampled are set to the values of the current best solution.

763 As for the binary case, we add all neighbors of the current best solution to the set of candidates and
764 we sample the 20 candidates with the highest acquisition function value.

765 We then use local search to optimize the acquisition function within the TR bounds while neighbors
766 are created by changing the index of one categorical variable. Again, we repeat until convergence
767 and return the best solution found during the local search.

768 **Ordinal problems.** The construction for ordinal problems is similar to the one for categorical
769 problems.

770 Suppose the number of ordinal variables of the problem is smaller or equal to the current TR length.
771 In that case, we sample an ordinal value uniformly at random to set the ordinal variable for each
772 candidate and each ordinal variable. Otherwise, we choose as many ordinal variables as each candi-
773 date’s current TR length and sample an ordinal value uniformly at random to set the ordinal variable.
774 We add all neighbors of the current best solution, all solutions where the distance to the current best
775 solution is 1 for one ordinal variable, to the set of candidates. We then sample the 20 candidates
776 with the highest acquisition function value and use local search to optimize the acquisition function
777 within the TR bounds. In the local search, increment or decrement the value of a single ordinal
778 variable.

779 **Mixed problems.** Mixed problems are effectively handled by treating every variable type sepa-
780 rately. Again, we create a set of initial random solutions where the values for the different variable
781 types are sampled according to the approaches described above. Note that this can lead to solutions
782 lying outside of the TR bounds. We simply remove these solutions and find the 20 best candidates
783 only across the solutions within the TR bounds.

784 When optimizing the acquisition function, we differentiate between continuous and combinatorial
785 variables. We optimize the continuous variables by gradient descent with the same settings as purely
786 continuous problems. When optimizing, we fix the values for the combinatorial values.

787 We use local search to optimize the acquisition function for the combinatorial variables. In this step,
788 we fix the values for the continuous variables and only optimize the combinatorial variables. We
789 create the neighbors by creating neighbors within Hamming distance of 1 for each combinatorial
790 variable type and then combining these neighbors. Again, we repeat the local search until conver-
791 gence.

792 We do five interleaved steps, starting with the continuous variables and ending with the combinatorial
793 variables.

794 **D Additional analysis of BODi and COMBO**

795 **D.1 Analysis of BODi**

796 **Binary problems.** BODi [18] is based on the idea of using a dictionary of reference points $\mathbf{A} =$
797 $(\mathbf{a}_1, \dots, \mathbf{a}_m)$ to encode a candidate point \mathbf{z} . In particular, the i -th entry of the m -dimensional
798 embedding $\phi_{\mathbf{A}}(\mathbf{z})$ is obtained by computing the Hamming-distance between \mathbf{z} and \mathbf{a}_i . Notably, the
799 dimensionality of the embedding m is chosen to be 128 in their experiments [18], which is larger
800 than the dimensionality of the benchmark functions themselves.

801 The dictionary elements \mathbf{a}_i are chosen such that they represent a wide range of *sequencies* where the
802 sequency of a binary string is defined as the number of times the string changes from 0 to 1 and vice
803 versa. Deshwal et al. [18] propose two approaches to generate the dictionary elements: (i) by using
804 binary wavelets, and (ii) by first drawing a Bernoulli parameter $\theta_i \sim \mathcal{U}(0, 1)$ for each $i \in [m]$ and
805 then drawing a binary string \mathbf{a}_i from the distribution $\mathcal{B}(\theta_i)$. The latter approach is their preferred
806 method.

807 We will now show that the probability of a point of sequence zero (i.e., $\mathbf{a}_i = \mathbf{0}$ or $\mathbf{a}_i = \mathbf{1}$) to
808 be sampled is higher than for arbitrary points. We hypothesize that BODi benefits from containing
809 such a point with high probability if the optimal point is also of sequency zero (cf. Section 4.6).
810 Since BODi’s performance degrades when randomizing the optimal point, we further hypothesize

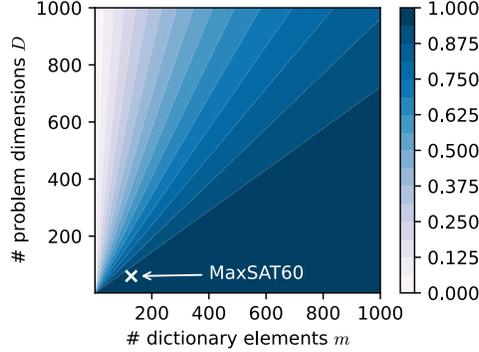


Figure 13: Probabilities of BODi to contain a zero-sequence solution for different choices of the dictionary size m and the function dimensionality D .

811 that BODi’s performance on problems with a zero- or low-sequence solution is not representative of
 812 problems with an arbitrary solution.

813 Deshwal et al. [18] choose to have 128 dictionary elements. Given a Bernoulli parameter θ_i , the
 814 probability that the i -th dictionary point \mathbf{a}_i is a point of sequency zero is given by $\theta_i^D + (1 - \theta_i)^D$:

$$\mathbb{P}(\text{“zero sequency”} \mid \theta_i) = \prod_{i=1}^m \theta_i^D + (1 - \theta_i)^D$$

815 Then, since θ_i follows a uniform distribution, the overall probability for a point of zero sequency is
 given by

$$\begin{aligned} \mathbb{P}(\text{“zero sequency”}) &= 1 - \underbrace{\prod_{i=1}^m \int_0^1 (1 - \theta_i^D - (1 - \theta_i)^D) \underbrace{p(\theta_i)}_{=1} d\theta_i}_{\text{prob. of } m \text{ times not zero sequency}} \\ &= 1 - \prod_{i=1}^m \left(\theta_i - \frac{\theta_i^{D+1}}{D+1} + \frac{(1 - \theta_i)^{D+1}}{D+1} \right) \Big|_{\theta_i=1} \\ &= 1 - \left(1 - \frac{2}{D+1} \right)^m, \end{aligned}$$

816 i.e., the probability of at least one dictionary element being of sequency zero is $1 - \left(1 - \frac{2}{D+1} \right)^m$.

817 The probability of BODi’s dictionary to contain a zero-sequence point increases with the number of
 818 dictionary elements m and decreases with the function dimensionality D (see Figure 13).

819 For instance, for the 60-dimensional MaxSAT60 benchmark, the probability that at least one dictio-
 820 nary element is of sequency zero is $1 - \left(1 - \frac{2}{60+1} \right)^{128} \approx 0.986$ (see Figure 13).

821 Note that there is at least one point \mathbf{z}^* with a probability of $\leq 1/2^d$ to be drawn. The probability of
 822 the dictionary containing that \mathbf{z}^* is less than or equal $1 - \left(1 - \frac{1}{2^d} \right)^m$ which is already less than 0.01
 823 for $d = 14$ and $m = 128$. In Section 4, we have shown that randomizing the optimal point structure
 824 leads to performance degradation for BODi. We hypothesize this is due to the reduced probability of
 825 the dictionary containing the optimal point after randomization.

826 **Categorical problems.** We calculate the probability that BODi contains a vector in its dictionary
 827 where all elements are the same. For categorical problems, BODi first samples a vector $\boldsymbol{\theta}$ from the
 828 τ_{\max} -simplex $\Delta^{\tau_{\max}}$ for each vector \mathbf{a}_i in the dictionary, with τ_{\max} being the maximum number of
 829 categories across all categorical variables of a problem. We assume that all variables have the same
 830 number of categories as is the case for the benchmarks in Deshwal et al. [18]. Let τ be the number
 831 of categories of the variables. For each element in \mathbf{a}_i , BODi then draws a value from the categorical

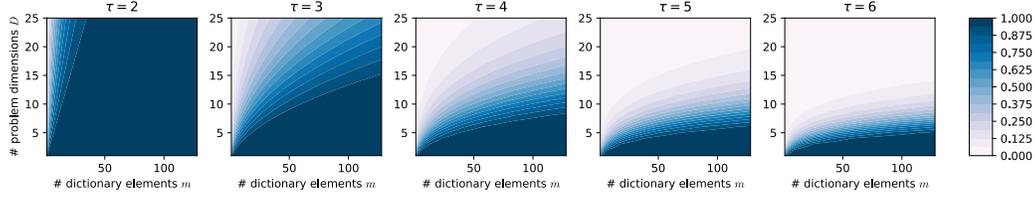


Figure 14: Probabilities of BODi’s dictionary to contain at least one categorical point where each category has the same value. The probability increases with the number of dictionary elements m but decreases with the number of categories τ and the number of problem dimensions D .

832 distribution with probabilities θ . While line 7 in Algorithm 5 in Deshwal et al. [18] might suggest
 833 that the elements in θ are shuffled for every element in \mathbf{a}_i , we observe that θ remains fixed based
 834 on the implementation provided by the authors⁴. The random resampling of elements from θ is
 835 probably only used for benchmarks where the number of realizations differs between categorical
 836 variables.

837 Then, for a fixed θ , the probability that all D elements in \mathbf{a}_i for any i are equal to some fixed value
 838 $t \in \{1, \dots, \tau\}$ is given by θ_t^D . The probability that, for any of the m dictionary elements, all D
 839 elements in \mathbf{a}_i are equal to some fixed value $t \in \{1, \dots, \tau\}$ is given by

$$\mathbb{P}(\text{“all one specific category”}) = 1 - \prod_{i=1}^m \int (1 - \theta_t^D) p(\theta_t) d\theta_t. \quad (3)$$

840 We note that θ follows a Dirichlet distribution with $\alpha = \mathbf{1}$ [85]. Then, θ_t is marginally Beta($1, \tau - 1$)-
 841 distributed [85]. With that, Eq. (3) becomes

$$\begin{aligned} \mathbb{P}(\text{“all one specific category”}) &= 1 - \prod_{i=1}^m \mathbb{E}_{\theta_t \sim \text{Beta}(1, \tau - 1)} [1 - \theta_t^D] \\ &= 1 - \prod_{i=1}^m (1 - \mathbb{E}_{\theta_t \sim \text{Beta}(1, \tau - 1)} [\theta_t^D]) \end{aligned}$$

842 now, by using the formula $\mathbb{E}[x^D] = \prod_{r=0}^{D-1} \frac{\alpha + r}{\alpha + \beta + r}$ for the D -th raw moment of a Beta(α, β) distri-
 bution [85]

$$\begin{aligned} &= 1 - \left(1 - \prod_{r=0}^{D-1} \frac{1 + r}{\tau + r} \right)^m \\ &= 1 - \left(1 - \frac{1}{\tau} \cdot \frac{2}{\tau + 1} \cdot \dots \cdot \frac{D}{\tau + D - 1} \right)^m \\ &= 1 - \left(1 - \frac{D! \tau!}{(\tau + D - 1)!} \right)^m \end{aligned}$$

843 We discussed in Section 4.3 that the PestControl benchmark obtains a good solution at $\mathbf{x} = \mathbf{5}$.
 844 One could assume that BODi performs well on this benchmark because its dictionary has a high
 845 probability to contain this point. However, we observe that the probability is effectively zero for
 846 $\tau = 5$, $m = 128$, and $D = 25$ (see Figure 14), which are the choices for the PestControl
 847 benchmark in Deshwal et al. [18]. This raises the question of (i) whether our hypothesis is wrong,
 848 and (ii) what the reason for BODi’s performance degradation on the PestControl benchmark is.

849 We show that BODi’s reference implementation differs from the algorithmic description in an impor-
 850 tant detail, causing BODi to be considerably more likely to sample category 5 on PestControl (or
 851 the “last” category for arbitrary benchmarks) than any other category.

⁴See https://github.com/aryandeshwal/bodi/blob/aa507d34a96407b647bf808375b5e162ddf10664/bodi/categorical_dictionary_kernel.py#L18

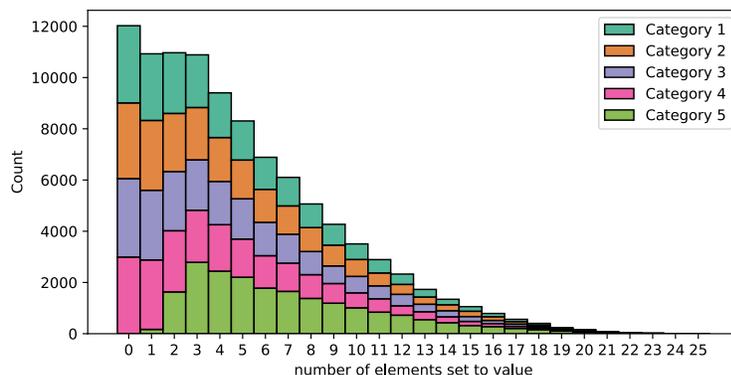


Figure 15: Histograms over the number of dictionary element entries set to each category for 20,000 repetitions of the sampling of dictionary elements for the `PestControl` benchmark. For each of the five categories and each value on the x -axis, the figure shows how often the number of entries in a dictionary element equals the value on the x -axis for the given category. For example, the count for $x = 0$ and category 5 is zero, indicating that all of the 20,000 dictionary points had at least one entry ‘5’. There is a considerably higher chance for a dictionary element entry to be set to category 5 than to any of the other categories.

852 In Figure 15, we show five histograms over the number of dictionary elements set to each category.
 853 The values on the x -axis give the number of elements in a 25-dimensional categorical vector being
 854 set to a specific category. One would expect that the histograms have a similar shape regardless of
 855 the category. However, for category 5, we see that more elements are set to this category than for the
 856 other categories: The probability of k elements being set to category 5 is almost twice as high as the
 857 probability of being set to another category for $k \geq 3$. In contrast, the probability that no element in
 858 the vector belongs to category 5 is virtually zero. This behavior is beneficial for the `PestControl`
 859 benchmark, which obtained the best value found during our experiments for $x^* = (5, 5, \dots, 5, 1)$
 860 (see Section 4). While we see that the probability of each dictionary entry being set to category 5 is
 861 very low, we assume that we sample sufficiently many dictionary elements within a small Hamming
 862 distance to the optimizer such that `BODi`’s GP is able to use this information to find the optimizer.

863 The reason for the oversampling of the last category lies in a rounding issue in the sampling of
 864 dictionary elements. In particular, for a given dictionary element \mathbf{a}_i and a corresponding vector $\boldsymbol{\theta}$
 865 with $|\boldsymbol{\theta}| = \tau$, for each $i \in \{1, \dots, \tau - 1\}$, Deshwal et al. [18] set $\lfloor D\boldsymbol{\theta}_i \rfloor$ elements to category i .
 866 The remaining $D - \sum_{i=1}^{\tau-1} \lfloor D\boldsymbol{\theta}_i \rfloor$ elements are then set to category τ . This causes the last category
 867 to be overrepresented in the dictionary elements. For the choices of the `PestControl` benchmark,
 868 $D = 25$ and $\tau = 5$, the first four categories had a probability of ≈ 0.1805 while the last one
 869 had a probability of ≈ 0.278 for 10^8 simulations⁵. We assume that the higher probability of the
 870 last category is the reason for the performance difference between the modified and the unmodified
 871 version of the `PestControl` benchmark.

872 D.2 COMBO on categorical problems

873 On the categorical `PestControl` benchmark, we could observe a similar behavior for `COMBO` [45]
 874 as for `BODi`.

⁵The 95% confidence intervals for categories 1–5 are (0.1799, 0.1807), (0.1802, 0.1810), (0.1803, 0.1811), (0.1801, 0.1809), (0.2775, 0.2783). Pairwise Wilcoxon signed-rank tests between categories 1–4 and category 5 gives p values of 0 ($W \approx 4.7 \cdot 10^{10}$ each).

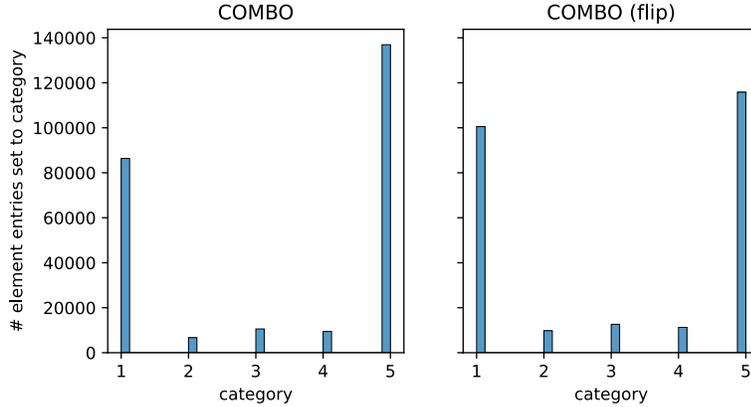


Figure 16: Histograms over the number of dictionary element entries set to each category for 20,000 repetitions of the sampling of dictionary elements for the `PestControl` benchmark. The histogram shows how many element entries of a 25-dimensional dictionary element are set to each of the five categories. There is a considerably higher chance for a dictionary element entry to be set to category 1 or 5 than to one of the other categories.

875 In Figure 16, we show the histograms over the number of dictionary elements set to each category
 876 for both the modified and the unmodified version of the `PestControl` benchmark. We see that the
 877 first and the last categories on both versions of the benchmark are overrepresented. As discussed
 878 in Section 4, this benchmark attains its best value for $x^* = (5, 5, \dots, 5, 1)$. Therefore, it seems
 879 unexpected that COMBO sets so many entries to category 1 for the unmodified benchmark version.
 880 For the modified benchmark version, this is entirely unexpected as the optimizer has a random
 881 structure. Here, one would expect the histogram to be uniform.

882 We argue that this behavior is at least partially caused by implementation error in the construction
 883 of the adjacency matrix and the Laplacian for categorical problems⁶. This error causes categorical
 884 variables to be modeled like ordinal variables. According to Oh et al. [45], categorical variables are
 885 modeled as a complete graph (see Figure 17).

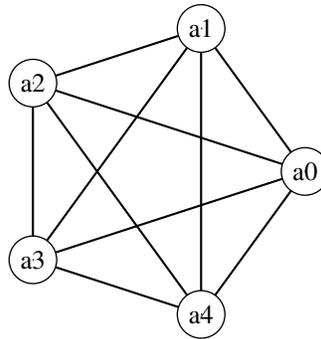


Figure 17: A categorical variable with five categories is modeled as a complete graph.

886 However, we find the adjacency matrix for the first category of a categorical variable with five
 887 categories is constructed as

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix},$$

⁶https://github.com/QUVA-Lab/COMBO/blob/9529eabb86365ce3a2ca44fff08291a09a853ca2/COMBO/experiments/test_functions/multiple_categorical.py#L137, last access: 2023-04-26

888 which is the adjacency matrix for a path graph with five vertices. We assume that the search space has
889 boundaries due to treating categorical variables as ordinal variables. Due to the high dimensionality
890 of the search space, COMBO visits the boundaries of the search space more often than the interior.

891 **E Extended related work**

892 Kim et al. [80] use a random projection matrix to optimize combinatorial problems in a continuous
893 embedded subspace. When evaluating a point, their approach first projects the continuous candidate
894 point to the high-dimensional search space and then rounds to the next feasible combinatorial solu-
895 tion. Deshwal et al. [71] build up on COMBO [45] by establishing a closed-form expression for the
896 diffusion kernel and proposing kernels tailored for mixed spaces.

897 **Monte-Carlo Tree Search.** Recent approaches employed Monte-Carlo Tree Search (MCTS) to
898 reduce the complexity of the problem. Wang et al. [95] use MCTS to learn a partitioning of the
899 continuous search space to focus the search on promising regions in the search space. Song et al.
900 [92] use a similar approach but instead of learning promising regions in the search space, they
901 assume an axis-aligned active subspace and use MCTS to select important variables.

902 **Non-linear embeddings.** Linear embeddings and random linear embeddings [37, 42, 48, 64, 70]
903 only require little or no training data to construct the embedding but assume a linear subspace. Non-
904 linear embeddings allow to learn more complex embeddings but often require more training data. Lu
905 et al. [82] and Maus et al. [83] use variational autoencoders (VAEs) to learn a non-linear embedding
906 of highly-structured input spaces. Tripp et al. [93] also use a VAE to learn a non-linear subspace of a
907 combinatorial search space. By using a re-weighting scheme that puts more emphasis on promising
908 points in the search space, they tailor the embedding toward optimization problems.

909 **Other combinatorial problems.** Deshwal et al. [72] propose two algorithms for *permutation*
910 *spaces* which occur in problems such as compiler optimization [29] and pose a special challenge
911 due to the superexponential explosion of solutions.

912 **References**

- 913 [5] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, An-
914 drew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo
915 Bayesian Optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*,
916 volume 33, 2020.
- 917 [6] Ricardo Baptista and Matthias Poloczek. Bayesian optimization of combinatorial structures.
918 In *International Conference on Machine Learning*, pages 462–471. PMLR, 2018.
- 919 [70] Mohamed Amine Bouhlel, Nathalie Bartoli, Rommel G. Regis, Abdelkader Otsmane, and
920 Joseph Morlier. Efficient global optimization for high-dimensional constrained problems by
921 using the Kriging models combined with the partial least squares method. *Engineering Opti-*
922 *mization*, 50(12):2038–2053, 2018.
- 923 [71] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Bayesian optimization over
924 hybrid spaces. In *International Conference on Machine Learning*, pages 2632–2643. PMLR,
925 2021.
- 926 [72] Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, and Dae Hyun Kim. Bayesian opti-
927 mization over permutation spaces. In *Proceedings of the AAAI Conference on Artificial Intelli-*
928 *gence*, volume 36, pages 6515–6523, 2022.
- 929 [18] Aryan Deshwal, Sebastian Ament, Maximilian Balandat, Eytan Bakshy, Janardhan Rao Doppa,
930 and David Eriksson. Bayesian Optimization over High-Dimensional Combinatorial Spaces
931 via Dictionary-based Embeddings. In *International Conference on Artificial Intelligence and*
932 *Statistics*, pages 7021–7039. PMLR, 2023.
- 933 [74] David Eriksson and Martin Jankowiak. High-dimensional Bayesian optimization with sparse
934 axis-aligned subspaces. In Cassio de Campos and Marloes H. Maathuis, editors, *Proceedings*
935 *of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Pro-*
936 *ceedings of Machine Learning Research*, pages 493–503. PMLR, 27–30 Jul 2021.

- 937 [20] David Eriksson and Matthias Poloczek. Scalable Constrained Bayesian Optimization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume
938 130 of *Proceedings of Machine Learning Research*, pages 730–738. PMLR, 13–15 Apr 2021.
939
- 940 [21] David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable
941 Global Optimization via Local Bayesian Optimization. In *Advances in Neural Information
942 Processing Systems (NeurIPS)*, volume 32, pages 5496–5507, 2019.
- 943 [77] Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wil-
944 son. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration.
945 In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.
- 946 [29] Erik Hellsten, Artur Souza, Johannes Lenfers, Rubens Lacouture, Olivia Hsu, Adel Ejjeh,
947 Fredrik Kjolstad, Michel Steuwer, Kunle Olukotun, and Luigi Nardi. BaCO: A Fast and
948 Portable Bayesian Compiler Optimization Framework. In *ACM International Conference on
949 Architectural Support for Programming Languages and Operating Systems*, 2023.
- 950 [32] Yingjie Hu, JianQiang Hu, Yifan Xu, Fengchun Wang, and Rong Zeng Cao. Contamination
951 control in food supply chain. In *Proceedings of the 2010 Winter Simulation Conference*, pages
952 2678–2681. IEEE, 2010.
- 953 [80] Jungtaek Kim, Seungjin Choi, and Minsu Cho. Combinatorial bayesian optimization with
954 random mapping functions to convex polytopes. In *Uncertainty in Artificial Intelligence*, pages
955 1001–1011. PMLR, 2022.
- 956 [37] Ben Letham, Roberto Calandra, Akshara Rai, and Eytan Bakshy. Re-Examining Linear Em-
957 beddings for High-Dimensional Bayesian Optimization. In *Advances in Neural Information
958 Processing Systems (NeurIPS)*, volume 33, pages 1546–1558, 2020.
- 959 [82] Xiaoyu Lu, Javier Gonzalez, Zhenwen Dai, and Neil D Lawrence. Structured variationally
960 auto-encoded optimization. In *International conference on machine learning*, pages 3267–
961 3275. PMLR, 2018.
- 962 [83] Natalie Maus, Haydn Thomas Jones, Juston Moore, Matt Kusner, John Bradshaw, and Jacob R.
963 Gardner. Local latent space bayesian optimization over structured inputs. In *Advances in
964 Neural Information Processing Systems (NeurIPS)*, volume 35, 2022.
- 965 [42] Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. A framework for Bayesian Op-
966 timization in Embedded Subspaces. In *Proceedings of the 36th International Conference on
967 Machine Learning*, volume 97 of *Proceedings of Machine Learning Research (PMLR)*, pages
968 4752–4761, 09–15 Jun 2019.
- 969 [85] Kai Wang Ng, Guo-Liang Tian, and Man-Lai Tang. Dirichlet and related distributions: Theory,
970 methods and applications. 2011.
- 971 [45] Changyong Oh, Jakub Tomczak, Efstratios Gavves, and Max Welling. Combinatorial Bayesian
972 Optimization using the Graph Cartesian Product. *Advances in Neural Information Processing
973 Systems (NeurIPS)*, 32, 2019.
- 974 [48] Leonard Papenmeier, Luigi Nardi, and Matthias Poloczek. Increasing the Scope as You Learn:
975 Adaptive Bayesian Optimization in Nested Subspaces. In *Advances in Neural Information
976 Processing Systems (NeurIPS)*, volume 35, 2022.
- 977 [88] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation
978 and approximate inference in deep generative models. In *International conference on machine
979 learning*, pages 1278–1286. PMLR, 2014.
- 980 [53] Binxin Ru, Ahsan Alvi, Vu Nguyen, Michael A Osborne, and Stephen Roberts. Bayesian
981 optimisation over multiple continuous and categorical inputs. In *International Conference on
982 Machine Learning*, pages 8276–8285. PMLR, 2020.
- 983 [90] Kenan Šehić, Alexandre Gramfort, Joseph Salmon, and Luigi Nardi. LassoBench: A High-
984 Dimensional Hyperparameter Optimization Benchmark Suite for Lasso. In *First Conference
985 on Automated Machine Learning (Main Track)*, 2022.

- 986 [91] Francisco J. Solis and Roger J-B. Wets. Minimization by Random Search Techniques. *Mathematics of Operations Research*, 6(1):19–30, 1981.
987
- 988 [92] Lei Song, Ke Xue, Xiaobin Huang, and Chao Qian. Monte Carlo Tree Search based Variable Selection for High Dimensional Bayesian Optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 35, 2022.
989
990
- 991 [93] Austin Tripp, Erik Daxberger, and José Miguel Hernández-Lobato. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:11259–11272, 2020.
992
993
- 994 [61] Xingchen Wan, Vu Nguyen, Huong Ha, Binxin Ru, Cong Lu, and Michael A. Osborne. Think Global and Act Local: Bayesian Optimisation over High-Dimensional Categorical and Mixed Search Spaces. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10663–10674. PMLR, 18–24 Jul 2021.
995
996
997
998
- 999 [95] Linnan Wang, Rodrigo Fonseca, and Yuandong Tian. Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:19511–19522, 2020.
1000
1001
- 1002 [64] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando de Freitas. Bayesian Optimization in a Billion Dimensions via Random Embeddings. *Journal of Artificial Intelligence Research (JAIR)*, 55:361–387, 2016.
1003
1004
- 1005 [66] James T Wilson, Riccardo Moriconi, Frank Hutter, and Marc Peter Deisenroth. The reparameterization trick for acquisition functions. *NeurIPS Workshop on Bayesian Optimization*, 2017.
1006