

1 We thank all reviewers for the time spent on their comprehensive reviews and we are happy that they enjoyed reading our  
2 paper. We are particularly glad to see reviewers unanimously agreeing that our work is a well-motivated and significant  
3 contribution to the NeurIPS community, providing a needed solution for a problem with immediate applications across  
4 a wide range of fields. In this rebuttal, we respond to each reviewer’s minor comments individually, as there are no  
5 comments shared across reviews.

#### 6 **Reviewer 1**

7 Thank you for the extremely positive review. We will fix the one typo reported.

#### 8 **Reviewer 2**

9 Thank you for the strongly positive review, arguing for acceptance based on strong performance across our compre-  
10 hensive and thorough experimental section. In light of your comments around our definitions of  $\lambda_m$  and  $\lambda_g$ , we will  
11 endeavour to improve clarity in the final version, emphasising that  $\lambda_m$  and  $\lambda_g$  are scalar quantities representing our  
12 kernel’s two hyper-parameters. These parameters tune the contribution of highly non-contiguous sub-sequences in our  
13 kernel calculations. We will also address your comments about combining the description of our proposed framework  
14 into a single location.

#### 15 **Reviewer 3**

16 Thank you for the predominately positive review, stressing the novelty and importance of our work to the NeurIPS  
17 community. We regret that minor clarity issues around our use of genetic algorithms prevented you from having as  
18 pleasant reading experience as our other reviewers.

19 No general guarantees exist for the convergence of genetic algorithms over string spaces and their derivation would  
20 be a substantial contribution in its own right. However, genetic algorithms were designed with string optimisation  
21 problems in mind and are widely used to explore string spaces (see Whitley et al. [1994]). Note that exact convergence  
22 of our genetic algorithm is not a requirement for effective BO as exact acquisition maximisation is rarely achieved even  
23 when performing BO over continuous spaces. All that is typically required is the identification of a string that scores  
24 reasonably highly according to our acquisition function.

25 We respond to your numbered comments as follows:

- 26 1. Rubanova et al. [2020] also consider Bayesian optimisation over strings. However, their approach, using  
27 ensembles of neural networks and generative models for acquisition function maximisation, has a smaller  
28 scope than our work, seeming suitable for only shorter strings of up to 100 characters and unconstrained string  
29 spaces. We will discuss this contemporaneous work (published at the start of this month) in our final version.
- 30 2. In Figure 5, we consider a gene optimisation problem over very long strings of 3672 base pairs and we use  
31 our locally-constrained genetic algorithm to explore this space of strings. It is unclear how a random forest  
32 (typically a classification or regression model) could be used to provide optimisation over discrete (never-mind  
33 syntactically constrained) spaces and this is beyond the scope of our work.
- 34 3. There is a stream of work combining convolution kernels and GPs for NLP applications (such as the cited  
35 Beck and Cohn [2017]), including kernels for modelling trees and graphs. Our method could definitely be  
36 extended to such kernels, providing efficient optimisation over additional types of discrete structures. As we  
37 mention in our Discussion section, we plan to investigate these avenues in future work.
- 38 4. We have presented numerical results for the efficiency of ours (and competitors) BO routines in Figure 4 of the  
39 main paper and in Figures 11-17 of our Appendix. We will better signpost this in the final version.

#### 40 **Reviewer 4**

41 Thank you for the extremely positive review and helpful minor comments that we will address to improve clarity in the  
42 final manuscript. We now deal with your concern that our kernel parameter  $\lambda_m$  is redundant after kernel normalisation.  
43 As our kernel (defined in line 105) contains the sum over all sub-sequences containing up to  $n$  characters, each kernel  
44 evaluation is a linear combination of  $\lambda_m$  raised to each power in  $[2, \dots, 2n]$ . As the coefficients of this linear combination  
45 depends on the two considered strings, a dependence on  $\lambda_m$  remains even after the normalisation of line 111. The  
46 choice of both  $\lambda_m$  and  $\lambda_g$  is important to the efficacy of the kernel and this is explicitly demonstrated for molecule  
47 prediction in Figure 18. In the final version, this figure will be brought into the main paper alongside further clarification  
48 of the role of the kernel parameters.