Thanks to all the reviewers for their constructive feedback, we respond to the major points below (other comments and suggestions will be applied to the final version of the paper).

• Lower bounds for the regret and C_T. To the best of our knowledge, There are no existing minimax optimal bounds (simultaneously optimal for both the regret and C_T) available even for the convex case (the bounds for the algorithms in prior works are not proven to be minimax optimal either). For the static regret, the $\mathcal{O}(\sqrt{T})$ lower bound on the regret from the Online Convex Optimization (OCO) literature applies here as well. However, for C_T, due to the trade-off between the regret and C_T, it is impossible to obtain lower bounds because obtaining better C_T bounds is always possible through incurring higher regret. For instance, one can always choose the action 0 which leads to $C_T < 0$. In the adversarial setting with $W = 1$, $\mathcal{O}(\sqrt{T})$ bounds for $R_{1,T}^{(A,S)}$ and C_T is the best so far in prior works in the convex setting and Algorithm 1 achieves the same bounds in the (non-convex) submodular framework.

• Issues with the plots. We will increase the font size of the plots, and give names to the algorithms in the final version of the paper. The decline in the plot for Algorithm 1 is due to the fact that once the dual variables get large enough ($\Theta(V)$), $V\nabla f_{t-1}(x_{t-1}^{(k)})$ and $\lambda_t^{(k)}\nabla g_{t-1}(v_{t-1}^{(k)})$ in the update rule of $v_t^{(k)}$ are of the same order, the algorithm becomes less aggressive in terms of utility maximization and it tries to balance its budget consumption and overall utility. This decline verifies our improved theoretical guarantee for C_T (compared to [16] and [17]) as well.

• Compare the algorithms with Meta-Frank-Wolfe. The primal update of the OSPHG and OLFW algorithms is the Meta-Frank-Wolfe algorithm for submodular maximization applied to the Lagrangian function and these two algorithms take into account the budget consumption as well. So, Meta-Frank-Wolfe algorithm can be viewed as a special case of these two algorithms with dual variable being set to zero, and it obtains a higher overall utility at the expense of further violating the budget constraint.

• Obstacles/challenges of applying prior works to our framework and the novelties/ideas in this work.

Limitations of the OSPHG algorithm. For the adversarial setting, the OSPHG algorithm obtains a $R_{W,T}^{(A,S)}$ bound of $O(\sqrt{WT})$ and a C_T bound of $O(W^{1/4}T^{3/4})$ that could be adapted to obtain expected $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(T^{3/4})$ bounds for $R_T^{(S,S)}$ and C_T respectively in the stochastic setting. However, in order to obtain *better* C_T bounds in both settings and obtain *any* high probability bounds in the stochastic setting, a different approach is needed.

Limitations of the OLFW algorithm. For the OLFW algorithm, the expected regret bound is sub-optimal (In our paper, our dual update is such that $\lambda_t^{(k)}$ is $\mathcal{F}_{t-1}$-measurable and $g_{t-1}(x^*)$ is independent of $\mathcal{F}_{t-1}$, where $\mathcal{F}_t = \{g_\tau\}_{\tau=0}^{t-1}$, which makes it possible to conclude $\mathbb{E}[\lambda_t^{(k)}g_{t-1}(x^*)] = \mathbb{E}\big[\lambda_t^{(k)}\mathbb{E}[g_{t-1}(x^*)]\big] \leq 0$, whereas this term is the dominating $\mathcal{O}(T^{3/4})$ term in the regret analysis of the OLFW algorithm). Moreover, there are no performance guarantees in the adversarial setting. In fact, the update rule for the dual variable in the OLFW algorithm is only reasonable when a good estimate of the constraint functions are available which is not the case in the adversarial setting.

How to deal with convex constraints? Neither of OSPHG and OLFW algorithms are able to deal with online convex constraints. Both these algorithms apply the Meta-Frank-Wolfe algorithm to the Lagrangian as the update rule for the primal variable. Although this approach makes it possible to get sub-linear $(1-\frac{1}{e})$-regret bounds, it is not the ideal way to treat the constraints and that is why they need to further restrict the constraint functions to be linear (i.e., ∇g_t being fixed) to obtain performance bounds. In order to remedy this issue, we treat the utility and the constraint in the Lagrangian function differently and we use $v_{t-1}^{(k)}$ (as opposed to $x_{t-1}^{(k)}$) as the argument for ∇g_{t-1} in the update rule $v_t^{(k)} = \mathcal{P}_{\mathcal{X}}\left(v_{t-1}^{(k)} + \frac{1}{2\alpha}\big(V\nabla f_{t-1}(x_{t-1}^{(k)}) - \lambda_t^{(k)}\nabla g_{t-1}(v_{t-1}^{(k)})\big)\right)$ which makes it possible to deal with convex constraints.

K dual variables needed. Using a single dual variable (which is done in [15], [16], [17] and [23]) is not enough to obtain the $\mathcal{O}(\sqrt{T})$ bounds for C_T simultaneously in both adversarial and stochastic settings and we instead maintain K dual variables which further complicates the analysis.

Comparison with [15] and [23]. Compared to [15], they only obtain regret and C_T bounds in the adversarial setting for the convex problem (no stochastic analysis). Also, compared to [23], their adversarial analysis for the convex problem is only done for the special case with window length $W = 1$ and in the stochastic setting, they only obtain bounds in expectation and they do not provide high probability performance guarantees (while Theorem 5 in our work provides high probability bounds). In summary, through our proposed algorithms and their analysis, we manage to address all the limitations of each of the prior works while maintaining their strengths and we provide a unified approach for all online submodular maximization problems with online convex constraints.

• Define x_t^* in line 172. x_t^* is any arbitrary action in the domain which satisfies the corresponding constraint $g_t(x_t^*) \leq 0$ and it does not need to be the instantaneous maximizer at round t. We will make this point clearer.

• Mention that T is known in advance. We are assuming that the horizon T is known in advance and the parameters of the proposed algorithms are in terms of T. However, had we not known T in advance, we could have used the well-known doubling trick to obtain the same regret and C_T bounds with slightly worse constants. We will specify this point in the final version of the paper.

• Explanations for the choice of DR-submodular functions as the objectives in the experiments. More explanations and motivations for the choice of functions are provided in Appendix B (and will be added to the final version of the paper).