Dear all reviewers: we highly appreciate your valuable comments and will reflect your comments in the revision.

[R1] Interpretation of the shared component: ["...wouldn't two brain areas that encode a particular stimulus, but not necessarily "communicate " with one another...exhibit a significant relationship? "] We agree with the reviewer that even if we find a significant relationship between two regions via dSCA, this does not guarantee that they are communicating directly. As has been noted for other methods that seek to detect inter-regional connectivity changes (e.g. Friston et al, Neuroimage 1997), there is always the possibility of a third region sharing information with both source and target regions. This is one of the fundamental problems for analysing observational data in general, and many methods have been proposed to detect causality (e.g., Shimizu et al., JMLR 2006). However, this caveat would be true whether using averaged stimulus representations (as in our paper), or whether full trial-to-trial variability (as suggested by the reviewer). The major advantage of using averaged representations is to allow for task-related information to be marginalized in the analysis, demixing components that encode different task variables. Despite this limitation, our method is an important starting point for subsequent interventional studies that more explicitly test task-related communication in a causal manner. We agree with the reviewer that "task-related communication" is a bit confusing, so we will change this to "task-related information sharing". We will also add these points in the Discussion.

[R1] Why is the "Full information" source matrix useful? The reviewer is correct: if the effects of different task parameters on neural populations are independent, the results obtained from full information matrix and marginalized matrix are indifferent. However, if their effects are not independent, marginalization by non-interesting task parameters may unintentionally diminish the information of the task parameter of interest. Note that although this is also the case when we use dPCA for a single brain region, this point was not discussed from this perspective. Thus, we view it as our novel theoretical contribution. We will modify this sentence in the revision.

[R2] Interpretability of the shared component. This is important. In the right figure, we visualize the contributions (absolute weights) to the first shared component for each neuron in the simulation analysis (Fig. 2 in the current paper; also see Fig. 1e). dSCA correctly identified neurons that contributed to (**a**) *Stimulus*- and (**b**) *Decision*-related communication. We emphasize that dSCA's linearity makes it easy. Although we only analyzed the simulation data due to insufficient space and time, we will add these results and results from real datasets to appendix.

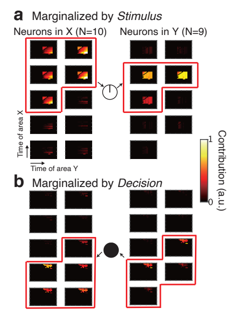

[R2] Novelty and originality. We insufficiently emphasized the novelty and originality of dSCA. Although several studies have investigated communication between populations of neurons [6-8], task-related communication has been ignored. This is of fundamental importance in neuroscience, and we show that it can be achieved simply by extending the previous method. While the reviewer argues it is a comparatively simple solution rather than an entirely new computational approach, we consider this a strength not a weakness because simplicity and ease of exposition are important points for practitioners. In sum, we believe that our approach is a novel, original, and useful.

[R2] Differentiation and advantages of dSCA. Both dPCA and NMF has been applied to obtain linear decompositions. The important contribution of dSCA is that it tests how task-relevant information is *shared* across brain regions, which is not the aim of dPCA or NMF. We will clarify this in the revision.

[R2] Correlations among task variables. We now realize that our explanation was insufficient. Although we only focused on marginalizing a single parameter (e.g. stimulus or decision) in the current paper, we can also marginalize for the interaction of multiple parameters, as was done in the dPCA paper [5]. We will clarify this in the revision.

[R2-R4] Assumptions and limitations. We agree that we should further describe the assumptions and limitations of dSCA. dSCA assumes that task-related communication is linearly represented. It makes dSCA simple and exactly solvable, and a linear method is popular in neuroscience because of its interpretability and less computational demand. However, this is also the limitation of dSCA: it cannot capture non-linear communication. There have been several methods for decomposing population neurons non-linearly, including deep learning. However, we are not aware of studies applied to neuroscience data for investigating task-related communication between multiple regions in low-dimensional projections of high-dimensional data. We believe our method is a good starting point for practitioners and methodological exploration. We will clarify these two points in the Methods and Discussion sections.

[R3] Regressing out before marginalization. If two task parameters are orthogonalized by experimental design, we do not need to do any procedure before marginalization (as in Fig. 4 in the current paper). However, if two task parameters are correlated by design (as in Fig. 3), to focus on a task parameter, we need to regress the other task parameters out from neural data before marginalization. We will clarify this point in the revision.

[R3] Scaling of explained variance We now realize that a more conventional definition of the explained variable is subtracting these values from one. We will change this in the revision.

[R1-R4] Other comments. Though we cannot address all comments, we assure reviewers we will do so in the revision.