We thank the reviewers for their detailed and insightful feedback! We are happy that the reviewers find our problem relevant and our evaluation and user study sound and significant. The reviewers recognize that we introduce a novel formulation of the assistance problem—tested for the first time with humans—with significant potential to spawn further future work in this direction. Due to space constraints, we focus our author response on addressing R1 and R4's core concern, comparing our proxy against true empowerment, and also provide some clarifications.

Empirical comparisons of the proxy with empowerment (R1, R3, R4): As it is not computationally feasible to directly compute empowerment in most domains, especially in real-time with humans in the loop, we motivate our proxy with experiments in domains where computing empowerment is feasible: 1) the known empowerment landscape of the non-linear pendulum 2) as suggested by R4, we evaluate the proxy method in the gridworld.

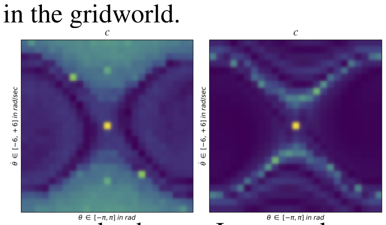

Following the intuition of the empowerment as quantifying the diversity of future states, our proxy estimates diversity by variance of final states. We find empirically in the well-known testbed of a non-linear pendulum that our proxy captures the essential properties of the empowerment landscape: the maximum is at the upright position of the pendulum (center of plot) and empowerment values are comparatively low for the states with energy below the separatrix energy. The left plot is the landscape of the proxy and the right plot is the corresponding empowerment landscape. Importantly, we do not aim to reproduce exact empowerment values, but rather to generally capture the analogous relative changes. The proxy indeed captures the critical qualitative features.

For the gridworld, replicating the experiments using the proxy instead of empowerment leads to a 1% decrease in success rate. Examining the failed cases shows that the proxy's inaccurate measure of empowerment can occasionally lead to blocking the goal, as in the non-ideal goal inference cases, however, the proxy failure rate is still much lower than that of goal inference (1% vs. 18%). These results suggest that the proxy appears to be a practical replacement for empowerment for the tested scenarios. Emphatically, the underlying intuition for our approach is to offer humans increasing controllability online; thus, computational efficiency is more critical than highly accurate empowerment.

Limitations of the proxy and of empowerment (R1, R4): We agree with the reviewers' suggestions that the paper would be strengthened with a more detailed treatment of the limitations of the proxy when used for assistance. We note that using the variance as a proxy for empowerment relies on an the assumption of homogeneous noise and high SNR, and would like to acknowledge R4's important insight that scenarios with noise varying between states would mislead our proxy, and will add these limitations to our discussion. To address R1's scalability comments, the proxy's sample-based method naturally requires exponentially more computational power as the action space grows. Nevertheless, there are many meaningful assistance applications with small action spaces (e.g. navigation with mobility devices, utensil stabilization). For large action spaces, we envisage that a natural way to achieve computational efficiency is by focusing on a subset of the action space most relevant for the range of tasks, which is an exciting area for future work. As for limitations with empowerment, as R4 suggests, predominantly optimizing for empowerment can lead to failure modes where the assistant prevents landing altogether (shown in the supplementary video); we will emphasize this point. R3 asks if there are environments that challenge/mislead empowerment. As our discussion acknowledges, we do not claim that empowerment alone is a general solution to all assistive tasks. An example of this is when the goal of the human requires them to lose empowerment to reach it, e.g. when the goal is inside a tunnel.

Novelty of empowerment-based assistance (R4, R2): We clarify: while prior work has used empowerment in agent rewards, we are proposing a formalism for assistance that is new relative to existing assistance formulations. As R3 mentions, although the components of our method are not novel, using them in this way is. R4 asks how our approach differs from [40], and we clarify that while [40] conceptually discusses the idea of an agent optimizing for human empowerment as it relates to Asimov's 3 laws of robotics, our paper concretely proposes that assistance in common shared workspaces and shared autonomy tasks can be cast as an empowerment problem, and evaluates that idea with real users. We certainly took inspiration from [40] in our formulation, and will amend the statement accordingly.

Combination of goal-agnostic and goal-oriented assistance (R4): To clarify, R4 is correct that allowing the agent to observe the recent human actions will lead the agent to implicitly try to infer the goal location. However, this does not rely on an explicit candidate set of goals, and the agent might only implicitly run inference over some lower-dimensional latent space (whatever increases success). The landing reward is necessary in the shared autonomy case as we discuss in the above limitations: solely focusing on empowerment can lead to failed assistance. The future work we discuss hopes to combine the strengths of goal-inference assistive methods with goal-agnostic assistance, as there are cases where the strengths of one can overcome the weaknesses of the other and vice versa.

Effect of c_{emp} (R1): In the Lander studies, we find that increasing c_{emp} generally makes the copilot more inclined to focus on stabilization, but too high c_{emp} causes the copilot to override the user and focus on hovering in the air as it prioritizes the high empowerment state over the reward of actually reaching the goal. We will add this to the discussion.