We thank all reviewers for their valuable comments and suggestions. We'll incorporate suggestions and clarifications in the revision. We first address a shared point (by Reviewer 1 and 2) and then respond to each reviewer respectively.

Further ablation about sub-graphs (R1 and R2). We have provided an ablation study about the main components of our method (e.g., multi-level attention and group attention) in the main paper. Regarding the detailed contributions of different types of sub-graphs, we further design three variants with different graph partitioning strategies: 1) "w/o relation awareness" combines $\boldsymbol{o}_{\mathrm{KG},2}$ (room objects) and $\boldsymbol{o}_{\mathrm{KG},3}$ (collected objects). 2) "w/o history awareness" combines $\boldsymbol{o}_{\mathrm{KG},4}$ with $\boldsymbol{o}_{\mathrm{KG},2}$ and $\boldsymbol{o}_{\mathrm{KG},3}$, respectively. 3) "w/o history" removes all historical information. Results in following figures indicate that the effect of different types of awareness varies with respect to the games. No simple conclusions can be made regarding which type of awareness contributes the most to the final performance (e.g., "w/o relation awareness" and "w/o history awareness" behave differently in "zork1" and "zork3"). However, considering them collectively and learning to balance their importance lead to the improved performance of our method.

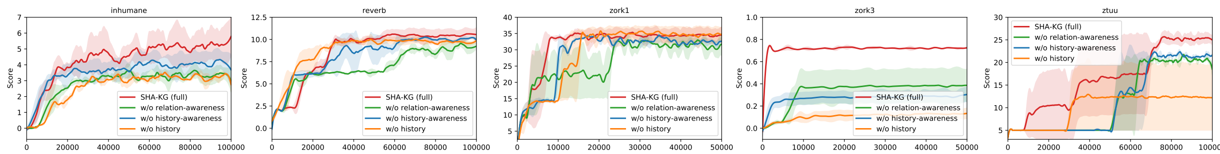

Reviewer 1 **Q1: Pre-defined rules. A1:** Our method does not have to rely on pre-defined rules to partition a knowledge graph. In fact, subgraphs are allowed to be constructed using different approaches. As there is no previous work considering subgraphs in text-based games, for simplicity, we use rules for graph partitioning, which can explicitly distinguish information and provide interpretability for the behaviors of the agent. In the future, we are interested in investigating other partitioning methods, such as language model-based question-answering (although the questions should be predefined), and automated partitioning.
Q2: What is missing from full KG that sub-graph captures? A2: The full KG can not distinguish between the current information and the historical information (note that SHA-KG also uses full KG in the high level). An example has been provided in Supplementary's Fig. 2 (b), where the yellow part denotes the current information and the green part indicates the historical information. From the yellow part, we can not tell whether an object has been collected (in inventory) or not. As shown in Fig. 2 (c), applying sub-graphs enables us to explicitly capture such information.

Reviewer 2 **Q1: Why SHA-KG architecture is leading to higher scores? A1:** Our two-level attention mechanism provides an explainable way to refine information. The first level of hierarchy tells the agent which part of the textual information should be focused on. Based on the output of this level, the second level of hierarchy informs the agent which part of a knowledge graph should be targeting.
Q2: Minor grammar mistakes and code. A2: We will address the grammar mistakes in the final version. Our code will be made publicly available upon publication.

Reviewer 3 **Q1: Controversy on attention. A1:** We will add discussion in the revision and cite the relevant papers.
Q2: What is the reasoning happening on KGs? A2: We will add an example to illustrate the graph partition and attention assignment, as suggested by the reviewer. Each graph indicates different types of information. In the right figure (based on Fig. 3 ztuu), for example, $\boldsymbol{o}_{\mathrm{KG},1}$ indicates room connectivity. We omit $\boldsymbol{o}_{\mathrm{KG},4}$ due to space limit. Our SHA helps the agent to use information efficiently for taking actions. The digit under the sub-graph denotes graph-level attention. Since the $\boldsymbol{o}_{\mathrm{KG},2}$ has the highest attention, the agent will focus more on objects it contains. In each sub-graph, nodes with top-3 highest attention (by GATs) are highlighted in yellow. Such node-level attention helps to further constrain (softly) the objects in $\boldsymbol{o}_{\mathrm{KG},2}$ to derive the actions.

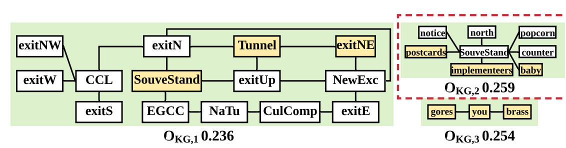

Q3: Graph mask for action selection. A3: All KG-related models construct masks from $\boldsymbol{o}_{\mathrm{KG,\,full}}$. As the benefit of graph mask has been investigated in KG-A2C, we use the same action selection strategy to make the comparison fair.

Reviewer 4 **Q1: Extensibility. A1:** Our method can be extended to handle continuous and realistic environments, if a knowledge graph can be constructed. For example, the entities of a knowledge graph may be extracted via object detection techniques for a visual environment. However, our hierarchical attention module is general to most tasks.
Q2: OpenIE. A2: We will add quantitative analysis for information extraction error.
Q3: Meaning of maximum score. A3: Yes, it refers to the maximum possible score in the game.
Q4: Multi-task setting. A4: The games used in this work have quite different characteristics, making it hard to share knowledge. Compared with Jericho, TextWorld may be a more suitable testbed for studying MTL, since it enables to generate a set of similar games. We leave the MTL as a future work.
Q5: Minor suggestions. A5: We will refine our paper based on these suggestions.