

A. Prompts and Responses

A.1. Human as an Observer

A.1.1. ROVER

Rover Domain: Prompt and GPT4 Response

Prompt:

Description: There is a rover in a 3x3 grid environment, working with a human manager. The rover is acting in the environment while the human manager is the observer in this case and can see the top-view.

Initial State: The rover starts at cell 5 and the human is observing the rover from the outside. There are resources to be collected which are represented as boxes in cells 2 and 4. There is one storage area that can store one resource which is represented as an open box in cell 3. The rover can also make observations but only after visiting that particular cell.

Actions: The rover actions include navigate from l_1 to l_2 , observe l_1 , load l_1 , and unload l_1 , each representing a set of actions since 1 (i.e., representing a location) can be instantiated to different cell locations, i.e., cells 0-8. Navigate (or nav) can move the rover from a location to one of its adjacent locations; load can be used to pick up a resource when the rover is not already loaded; unload can be used to unload a resource at a storage area if the area is empty; observe (or obs) can be used to make an observation. Once a location is observed, it remains observed.

Goal: The goal is for the rover to make the storage area non-empty in cell 3 and observe two locations in cells 0 and 6. But the rover does not have knowledge of these cells, and hence, needs to keep observing all cells till it finds the correct ones.

Definitions: Explicable actions: any agent (rover) action that is easily understood by the observer with respect to its contribution in achieving the final task will be called explicable actions.

Rover’s plan: starts from cell 5, navigates to cell 4, loads a resource, navigates to cell 1, navigates back to cell 4, navigates to cell 3, navigates to cell 0, observes the cell, navigates to cell 3, unloads the resource, navigates to cell 6, observes the cell, and terminates the task.

Question: Imagine you are the human in this case looking at the rover executing the above plan. Would you find the rover’s actions of navigating to cell 1 from cell 4 and then navigating back to cell 4 explicable or not? Give reason.

GPT4 Response:

As a human observer in this case, I would not find the rover’s actions of navigating to cell 1 from cell 4 and then navigating back to cell 4 to be explicable. The reason is that these actions do not seem to contribute to achieving the final task directly. Instead, they appear to be unnecessary movements that do not help in loading or unloading the resource, or in observing the required locations.

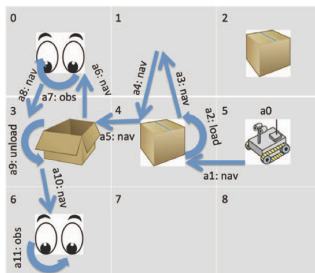


Figure 2: Rover domain, adapted as in (Zhang et al., 2017).

Human Ground Truth: Human does not understand the rover oscillating between two cells. However, the rover does it to observe the adjoining cell and come back in step 4 as the goal is to observe for locations represented with an eye.

A.1.1.2. FETCH

Fetch Domain: Prompt and GPT4 Response

Prompt:

Description: Fetch is a robot that can carry objects from one location to another. Fetch robot’s design requires it to tuck its arms and lower its torso or crouch before moving - which is not obvious to a human navigating it. This may lead to an unbalanced base and toppling of the robot if the human deems such actions as unnecessary.

Now, consider a problem with the following initial and goal states (here, identical for both the robot and the human)

- **Initial state:** There is a block b1 at location loc1, and the robot is at location loc1 and has its hand empty:

```
(:init (block-at b1 loc1) (robot-at loc1) (hand-empty))
```

Goal: The robot has to pick the block b1 and take it to location loc2:

```
(:goal (and (block-at b1 loc2)))
```

Definition: Explicable action: any agent (rover) action that is easily understood by the observer with respect to its contribution in achieving the final task will be called explicable actions.

Question: In the robot’s optimal plan, it will tuck its arms before moving from location 1 to location 2. Imagine you are the human in this case. Would you find such an action of the robot explicable or not? Give reason. —

GPT4 Response:

As a human observer, I might not find the robot’s action of tucking its arms and lowering its torso before moving from location 1 to location 2 explicable initially. This is because the crouching action might not seem necessary or directly contributing to the final task of picking up the block and moving it to location 2.

However, if I were provided with more information about the robot’s design requirements, and why the crouching action is crucial for its functionality and stability, I might eventually understand the necessity of the action and regard it as explicable.



Figure 3: Fetch domain, adapted as in (Chakraborti et al., 2017).

Human Ground Truth: An optimal plan for the robot, in this case, involves a tuck action followed by a move - pick-up b1 -> tuck -> move loc1 loc2 -> put-down b1

The human, on the other hand, expects a much simpler model, as shown below. The move action does not have the preconditions for tucking the arm and lowering the torso, while tuck does not automatically lower the torso either. Hence, the robot’s actions are inexplicable to the human.

A.1.3. URBAN SEARCH AND RESCUE (USAR)

USAR Domain: Prompt and GPT4 Response

Prompt:

Description: In a typical Urban Search and Rescue (USAR) setting, there is a building with interconnected rooms and hallways, with a human commander CommX (who has a top view of the building) and a robot which is acting in the environment. All of these agents are autonomous agents working together or independently in the same environment. **Actions:** The commander has capabilities to move and conduct triage at specified locations, and he can also meet with other agents, as well as pickup, drop off or handover medkits to accomplish their task. The robot can similarly move about, search rooms, or handover or change the position of the medkits. It can thus have its own goals (maybe from being directly assigned by the commander himself or due to long term task specifications), but can also help the commander in accomplishing his goals by fetching the medkits for him. **Initial State:** Suppose there are two medkits: medkit1 which is a heavier medkit that lies closer to the room where the human is, and medkit2 which is the lighter medkit that lies across the hall close to the room which is also the triage location. **Question:** Imagine you are the human observing this scenario from the top-view and do not have any other information about the medkits. Which medkit would you pick to carry to the triage location? —

GPT4 Response:

If I were the human observing this scenario from the top-view without any other information about the medkits, I would pick medkit2 because it is closer to the triage location and lighter. This would likely require less effort to transport to the triage location and be more efficient overall.

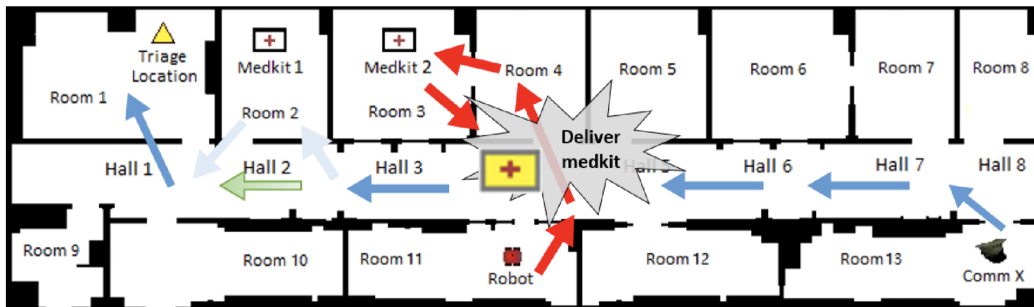


Figure 4: USAR domain, adapted as in (Chakraborti et al., 2015; Sreedharan et al., 2017).

Human Ground Truth: Human would pick the nearest medkit, i.e., medkit1 as they do not know which medkit is lighter or heavier and would want to get the one closest to them.

A.2. Human as an Actor

A.2.1. OVERCOOKED

Overcooked Domain: Prompt and GPT4 Response

Prompt:

Overcooked is based on a popular two-player video game in which players control chefs in a kitchen to cook and serve dishes. Each dish takes several high-level actions to deliver, making strategy coordination difficult, in addition to the challenge of motion coordination.

The domain has objects including onions, dishes, and soups. Players place 3 onions in a pot, leave them to cook for 20 time steps, put the resulting soup in a dish, and serve it, giving all players a reward of 20. The six possible actions are: up, down, left, right, noop, and interact, which does something based on the tile the player is facing, e.g. placing an onion on a counter.

Each layout has one or more onion dispensers and dish dispensers, which provide an unlimited supply of onions and dishes respectively. Agents should learn how to navigate the map (light brown area), interact with objects (such as ingredients, plates, and cooking stove on the counter shown in dark brown), drop the objects off in the right locations (such as putting onions to the cooking stove, cooked soup from the stove in the white plates, etc.), and finally serve completed dishes to the serving area (represented by the gray box). All the while, agents should be aware of what their partner is doing and coordinate with them effectively.

Imagine that you are a human agent and have been paired with an AI agent to play the game of Overcooked, where both you and the AI agent are going to accomplish a series of sub-tasks to accomplish the final goal of delivering as many cooked dishes as possible. List out 7 events (or sub tasks), you believe would be reasonable in this game to accomplish. Assume, that both you and AI agent have the same capabilities in this game, and thus can do any and all sub-tasks possible. There is no explicit mode of communication between you and the AI agent. —

GPT4 Response:

1. Retrieving onions: One of the agents (either human or AI) should approach the onion dispenser to collect onions.
2. Chopping onions: After collecting the onions, the agent should move to a chopping board and use the 'interact' action to chop the onions, preparing them for the soup.
3. Placing chopped onions in the pot: The agent with the chopped onions should navigate to the cooking pot and use the 'interact' action to place the onions inside.
4. Cooking the soup: Once three chopped onions are placed in the pot, agents should wait for 20 time steps for the soup to cook.
5. Collecting dishes: Meanwhile, the other agent (human or AI) should approach the dish dispenser to collect a dish.
6. Plating the cooked soup: After the soup is cooked, one of the agents should move to the pot, use the 'interact' action to collect the soup, and then transfer it to the collected dish.
7. Serving the finished dish: Finally, the agent carrying the plated soup should navigate to the serving area (gray box) and use the 'interact' action to serve the soup, earning a reward for both players.

Throughout these events, both the human and AI agents will need to be aware of each other's movements and actions to coordinate effectively and avoid blocking each other's paths or disrupting each other's tasks.

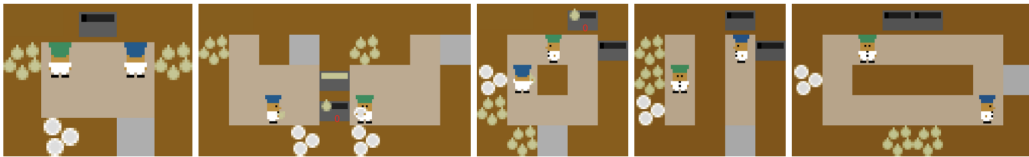


Figure 5: Overcooked domain, adapted as in (Carroll et al., 2019). Layouts from left to right: *Cramped Room*, *Asymmetric Advantages*, *Coordination Ring*, *Forced Coordination*, and *Counter Circuit*.

A.2.2. OVERCOOKED: LAYOUT 1 - ASYMMETRIC ADVANTAGES

Overcooked Domain: Asymmetric Advantages: Prompt and GPT4 Response

Prompt:

Overcooked is based on a popular two-player video game in which players control chefs in a kitchen to cook and serve dishes. Each dish takes several high-level actions to deliver, making strategy coordination difficult, in addition to the challenge of motion coordination.

The domain has objects including onions, dishes, and soups. Players place 3 onions in a pot, leave them to cook for 20 time steps, put the resulting soup in a dish, and serve it, giving all players a reward of 20. The six possible actions are: up, down, left, right, noop, and interact, which does something based on the tile the player is facing, e.g. placing an onion on a counter.

Each layout has one or more onion dispensers and dish dispensers, which provide an unlimited supply of onions and dishes respectively. Agents should learn how to navigate the map (light brown area), interact with objects (such as ingredients, plates, and cooking stove on the counter shown in dark brown), drop the objects off in the right locations (such as putting onions to the cooking stove, cooked soup from the stove in the white plates, etc.), and finally serve completed dishes to the serving area (represented by the gray box). All the while, agents should be aware of what their partner is doing and coordinate with them effectively.

We are in the Asymmetric Advantages layout which tests whether players can choose high-level strategies that play to their strengths. There is a counter in the middle with two stoves that can be accessed from each side. Both players have onion dispenser, plate dispenser and serving area on their sides. However, the plates and the serving area are closer to the player on the left, while the onion dispenser is closer to the player on the right.

Imagine that you are a human agent on the left and have been paired with an AI agent on the right, to play the game of Overcooked, where both you and the AI agent are going to accomplish a series of sub-tasks to accomplish the final goal of delivering as many cooked dishes as possible. List out 7 events (or sub tasks), you believe would be reasonable in this game to accomplish. Assume, that both you and AI agent have the same capabilities in this game, and thus can do any and all sub-tasks possible. There is no explicit mode of communication between you and the AI agent. —

GPT4 Response:

1. Event 1: Both players start at their respective positions in the kitchen. The AI agent moves towards the onion dispenser on their side, while you move towards your plate dispenser.
2. Event 2: The AI agent picks up an onion from their onion dispenser and places it onto the central counter. Simultaneously, you pick up a plate from your plate dispenser and place it onto the counter on your side.
3. Event 3: The AI agent collects two more onions from their dispenser and places them onto the counter next to the first onion, completing the necessary ingredients for a soup. Meanwhile, you stand near the stove, ready to start cooking.
4. Event 4: You pick up the three onions placed by the AI agent and put them into a pot on the stove. You then wait for the soup to cook for 20 time steps.
5. Event 5: While the soup is cooking, the AI agent collects more onions and places them onto the counter, preparing for the next dish. Meanwhile, you stand ready to collect the cooked soup from the stove.
6. Event 6: After the soup has cooked for 20 time steps, you take the pot off the stove and transfer the soup onto the plate you collected earlier. You then place the pot back onto the stove, ready for the next batch of onions.
7. Event 7: You carry the plated soup to the serving area and deliver it, achieving a reward of 20 for both players. Simultaneously, the AI agent begins the process again by placing the prepped onions into the now-empty pot on the stove. You both continue this cycle, trying to deliver as many dishes as possible within the given time.

A.2.3. OVERCOOKED: LAYOUT 2 - FORCED COORDINATION

Overcooked Domain: Forced Coordination: Prompt and GPT4 Response

Prompt:

Overcooked is based on a popular two-player video game in which players control chefs in a kitchen to cook and serve dishes. Each dish takes several high-level actions to deliver, making strategy coordination difficult, in addition to the challenge of motion coordination.

The domain has objects including onions, dishes, and soups. Players place 3 onions in a pot, leave them to cook for 20 time steps, put the resulting soup in a dish, and serve it, giving all players a reward of 20. The six possible actions are: up, down, left, right, noop, and interact, which does something based on the tile the player is facing, e.g. placing an onion on a counter.

Each layout has one or more onion dispensers and dish dispensers, which provide an unlimited supply of onions and dishes respectively. Agents should learn how to navigate the map (light brown area), interact with objects (such as ingredients, plates, and cooking stove on the counter shown in dark brown), drop the objects off in the right locations (such as putting onions to the cooking stove, cooked soup from the stove in the white plates, etc.), and finally serve completed dishes to the serving area (represented by the gray box). All the while, agents should be aware of what their partner is doing and coordinate with them effectively.

We are in the Forced Coordination layout which forces players to develop a high-level joint strategy, since neither player can serve a dish by themselves due to a counter table between them over which the player on the left side can pass over onions and plates, and the right player will take the onions, put them on the cooking stove, plate the cooked soup in a dish, and finally serve them.

Imagine that you are a human agent on the left side and have been paired with an AI agent to play the game of Overcooked, where both you and the AI agent are going to accomplish a series of sub-tasks to accomplish the final goal of delivering as many cooked dishes as possible. List out 7 events (or sub tasks), you believe would be reasonable in this game to accomplish. Assume, that both you and AI agent have the same capabilities in this game, and thus can do any and all sub-tasks possible. There is no explicit mode of communication between you and the AI agent. —

GPT4 Response:

1. Event 1: The AI agent collects onions from the onion dispenser while I collect dishes from the dish dispenser.
2. Event 2: The AI agent places the onions on the counter table, while I place dishes on the counter table. We both then collect more onions and dishes to prepare for the next round.
3. Event 3: I pick up the onions placed by the AI agent on the counter table and pass them over to the AI agent's side, while the AI agent picks up the dishes I placed on the counter table.
4. Event 4: The AI agent puts the onions on the cooking stove to cook the soup while I continue to collect more dishes and onions, passing them again to the AI agent.
5. Event 5: The AI agent monitors the cooking progress and takes the cooked soup out of the stove, placing it onto the dishes collected earlier.
6. Event 6: The AI agent moves the completed dishes with soup to the serving area, while I continue to prepare more onions and dishes for the next round.
7. Event 7: We both go back to our respective stations, The AI agent continues cooking soups and serving them, while I keep passing onions and dishes to the AI agent. Throughout the game, we efficiently coordinate our actions and maintain a continuous flow of cooked dishes being served.

A.2.4. OVERCOOKED: LAYOUT 3 - COUNTER CIRCUIT

Overcooked Domain: Counter Circuit: Prompt and GPT4 Response

Prompt:

Overcooked is based on a popular two-player video game in which players control chefs in a kitchen to cook and serve dishes. Each dish takes several high-level actions to deliver, making strategy coordination difficult, in addition to the challenge of motion coordination.

The domain has objects including onions, dishes, and soups. Players place 3 onions in a pot, leave them to cook for 20 time steps, put the resulting soup in a dish, and serve it, giving all players a reward of 20. The six possible actions are: up, down, left, right, noop, and interact, which does something based on the tile the player is facing, e.g. placing an onion on a counter.

Each layout has one or more onion dispensers and dish dispensers, which provide an unlimited supply of onions and dishes respectively. Agents should learn how to navigate the map (light brown area), interact with objects (such as ingredients, plates, and cooking stove on the counter shown in dark brown), drop the objects off in the right locations (such as putting onions to the cooking stove, cooked soup from the stove in the white plates, etc.), and finally serve completed dishes to the serving area (represented by the gray box). All the while, agents should be aware of what their partner is doing and coordinate with them effectively.

We are in the Counter Circuit layout which involves a non-obvious coordination strategy, where onions are passed over a counter in the middle of the kitchen to the pot, rather than being carried around the counter. There is only one path around the counter so the two agents can not cross each other and will collide if they reach the same location in the kitchen.

Imagine that you are a human agent on the left and have been paired with an AI agent on the right, to play the game of Overcooked, where both you and the AI agent are going to accomplish a series of sub-tasks to accomplish the final goal of delivering as many cooked dishes as possible. List out 7 events (or sub tasks), you believe would be reasonable in this game to accomplish. Assume, that both you and AI agent have the same capabilities in this game, and thus can do any and all sub-tasks possible. There is no explicit mode of communication between you and the AI agent. —

GPT4 Response:

1. Navigate to the onion dispenser: You move to the onion dispenser on the left side of the kitchen, while the AI agent waits near the counter.
2. Collect onions: You pick up an onion from the dispenser using the "interact" action and bring it to the counter in the middle of the kitchen.
3. Pass onions to the AI agent: You place the onion on the counter, and the AI agent picks it up and moves to the pot.
4. AI agent places onions in the pot: The AI agent puts the onion in the pot using the "interact" action. This process is repeated until there are three onions in the pot.
5. AI agent collects dishes: While the soup is cooking, the AI agent moves to the dish dispenser to grab a dish, preparing it for the cooked soup.
6. AI agent transfers cooked soup to the dish: After waiting for the soup to cook, the AI agent picks up the cooked soup using the "interact" action, placing it on the collected dish.
7. AI agent serves the completed dish: Finally, the AI agent moves to the serving area and delivers the completed dish using the "interact" action. This process is then repeated to maximize the number of dishes served.

Ground Truth Event List:

- putting an onion/dish/soup on the counter,
- picking up an onion/dish/soup from the counter,
- picking up an onion from the onion dispenser,
- picking up a dish from the dish dispenser,
- picking up a ready soup from the pot with a dish,

- placing an onion into the pot,
- delivering a soup to the serving area.

B. User Study

We conduct a study with 17 participants, and the study consists of 2 parts. Part 1, as shown in Fig. 6 provides a domain description and the task details to the users. This also helps in familiarizing the users with the domain. Note, that this is the same as the prompt we use for querying all the LLMs. The users are also provided with additional details on the usage of the study interface. Once they add events to complete Task 1, they move to Part 2 which is shown in Fig. 7. Part 2 presents another list of events alongside the list of events the user submitted in Part 1. This other list (List A in Fig. 7(a)) is the response generated by GPT-4. Finally the specific questions asked to the user for the Part 2 is shown in Fig. 7(b).

The study has IRB approval and ensured that participant confidentiality, informed consent, and voluntary participation is taken in accordance with the IRB guidelines.

B.1. Interface

User Study : Part 1

Overcooked is based on a popular two-player video game in which players control chefs in a kitchen to cook and serve dishes. Each dish takes several high-level actions to deliver, making strategy coordination difficult, in addition to the challenge of motion coordination.

The domain has objects including onions, dishes, and soups. Players place 3 onions in a pot, leave them to cook for 20 time steps, put the resulting soup in a dish, and serve it, giving all players a reward of 20. The six possible actions are: up, down, left, right, noop, and interact, which does something based on the tile the player is facing, e.g. placing an onion on a counter.

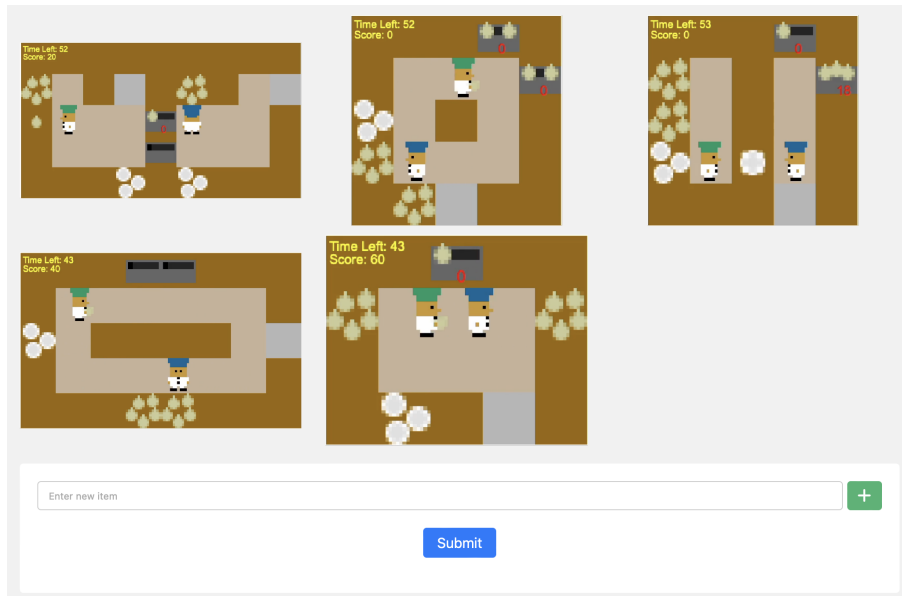
Each layout has one or more onion dispensers and dish dispensers, which provide an unlimited supply of onions and dishes respectively. Agents should learn how to navigate the map (light brown area), interact with objects (such as ingredients, plates, and cooking stove on the counter shown in dark brown), drop the objects off in the right locations (such as putting onions to the cooking stove, cooked soup from the stove in the white plates, etc.), and finally serve completed dishes to the serving area (represented by the gray box). All the while, agents should be aware of what their partner is doing and coordinate with them effectively.

Your Objective: Imagine that you have been paired with an AI agent to play the game of Overcooked, where both you and the AI agent are going to accomplish a series of sub-tasks to accomplish the final goal of delivering as many cooked dishes as possible. List out 7 events (or sub tasks), you believe would be reasonable in this game to accomplish. Assume that both you and the AI agent have the same capabilities in this game, and thus can do any and all sub-tasks possible. There is no explicit mode of communication between you and the AI agent.

Instructions :

1. Write each event in the box given below, and press "Enter" on the keyboard or click on '+' icon on the right.
2. You can also delete any event you add by clicking on the 'bin' icon on the right.
3. The events do not need to be in any particular order.
4. Please ensure you do not repeat the events by paraphrasing.

(a) Part 1 - User Study Description



(b) Part 1 - User Study Game Visualization

Figure 6: User Study: Part 1

User Study : Part 2

<p>List A : (unknown source)</p> <p>Collect onions: I will collect onions from the onion dispenser and place them onto the nearest available counter.</p> <p>Chop onions: When there are onions on the counter, I will chop the onions and place the chopped onions on a separate counter, ready to be added to the pot.</p> <p>Add onions to the pot: When there are three chopped onions, I will place all three of them in a pot to start cooking the soup.</p> <p>Collect dishes: I will collect dishes from the dish dispenser and place them on a nearby counter, ready for the soup to be served.</p> <p>Transfer cooked soup to dishes: Once the soup is cooked for 20 time steps, I will transfer the cooked soup to the empty dishes on the counter.</p> <p>Deliver completed dishes: After transferring the cooked soup to a dish, I will carry the completed dish and deliver it to the serving area.</p> <p>Clean and reset: If necessary, I will clean up any dropped onions or dishes to maintain a tidy kitchen. Moreover, I will reset the cooking station to make it available for the next round of ingredients.</p>	<p>List B : (Created by you in the previous step)</p> <p>Event 1</p> <p>Event 2</p> <p>Event 3</p> <p>Event 4</p> <p>Event 5</p> <p>Event 6</p> <p>Event 7</p>
--	---

(a) Part 2 - Event Lists

Below you can find two lists that also gives the events. List A is from an unknown source and List B was given by you in the previous step. Evaluate List A by answering the following questions on a scale of 1 to 5, where 5 represents complete agreement and 1 represents complete disagreement.

How well do you think the List A is aligned to your event list?
_____○_____ 3

Considering the content of the listed events, could you determine if they were generated by a human user expressing their preferences?
_____○_____ 3

Submit

(b) Part 2 - Questions

Figure 7: User Study: Part 1