

A Reflexion

A.1 Architecture Description using FALAA

The following sections present the information retrieval and prompt creation fragments of the **Planner** and **Reflector** components of the **Reflexion** agent in Figures 10 and 11, respectively.

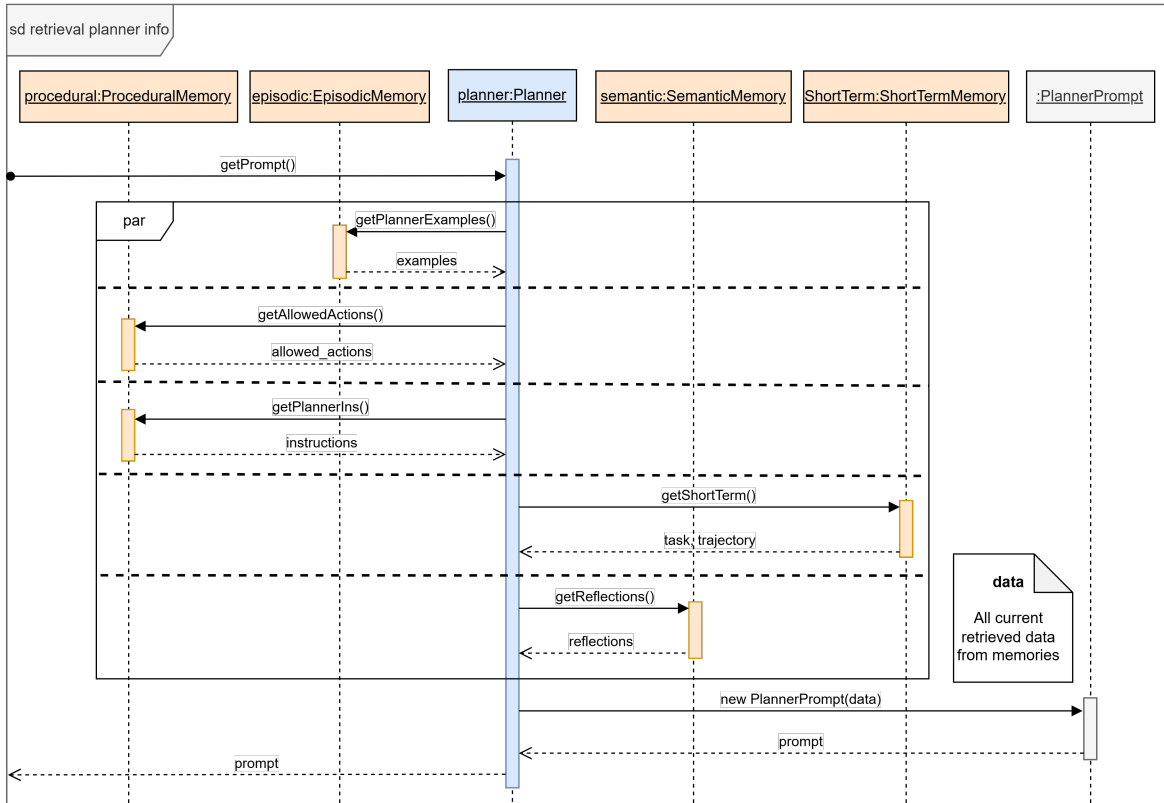


Fig. 10: UML sequence diagram of the information retrieval process for **Planner** to generate a prompt. In this flow, **Planner** uses internal actions of type **retrieval**, represented by the *get* methods of the **episodic**, **procedural**, and **semantic** memories, from which examples are obtained along with the allowed actions and instructions. These are condensed into the *PlannerPrompt* object, aiming to guide the behavior of the **Planner** LLM to generate actions.

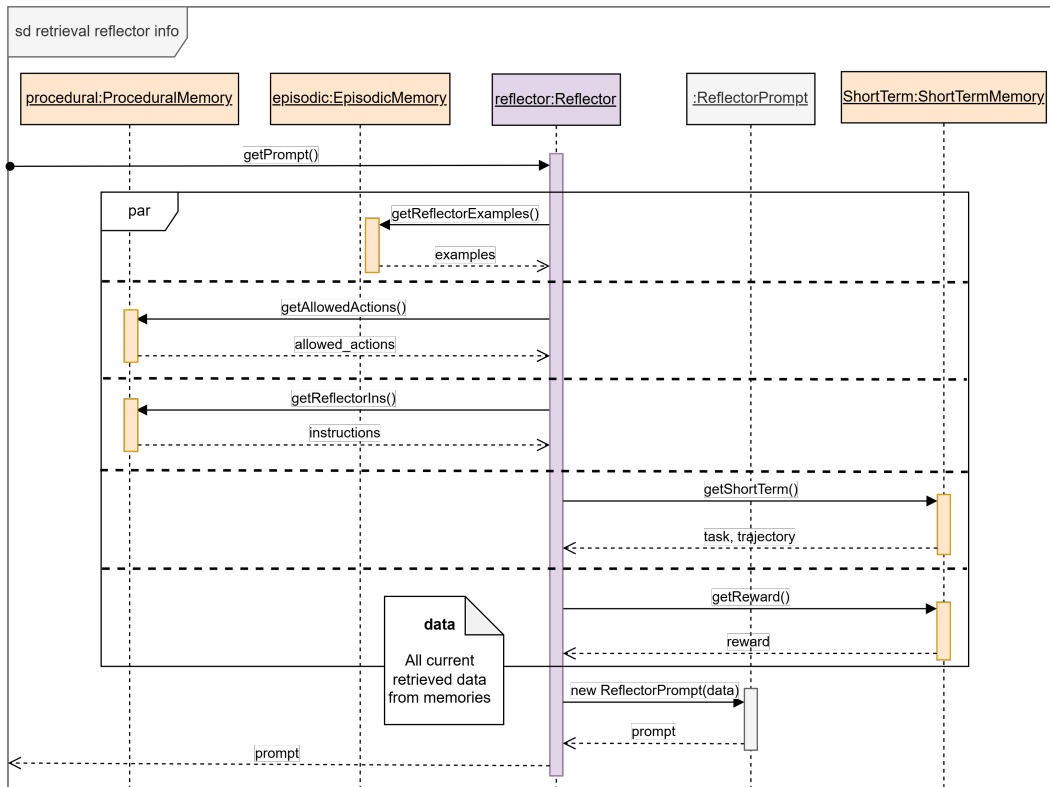


Fig. 11: UML sequence diagram of the information retrieval process for **Reflector** to generate a prompt. In this flow, **Reflector** uses internal actions of type **retrieval**, represented by the *get* methods of the **episodic** and **procedural** memories, from which examples for the LLM and the allowed actions along with the instructions to guide the LLM behavior are obtained. These, together with the task and current trajectory retrieved from the **short-term** memory, are used to generate the prompt.

B Retroformer

B.1 Pseudo Code

Figure 12 shows the training process based on RLHF, which was ultimately used as a guide to decide whether **Retroformer** would be trained for one or multiple environments.

Algorithm 1 **Retroformer** with Policy Gradient Optimization

- 1: Initialize TEXT-DAVINCI-003 as the Retrospective model with LONGCHAT-16K. Set the maximum trials for rollouts as $N = 3$. The temperature used for sampling $t_s = 0.9$.
 - 2: **Step 1: Offline Data Collection.** Collect multiple rollouts for each environments k ($k = 1, \dots, K$) for the tasks in the training sets and save as D_{RL} .
 - 3: **for** episode $t = 1, \dots, N$ **do**
 - 4: **for** source domain $k = 1, \dots, K$ **do**
 - 5: Receive trajectory $[s_{k,i,\tau}, a_{k,i,\tau}, r_{k,i,\tau}]_{\tau=1}^T$ and episodic returns $G_{k,i}$ for task i .
 - 6: **for** unsuccessful tasks j **do**
 - 7: Randomly sample a pair of reflection responses $(y_{k,j}^{(1)}, y_{k,j}^{(2)})$ with Retrospective LM temperature set to t_s , with the same instruction prompt defined in Eq. (4).
 - 8: Roll out the next episode with $y_{k,j}$, and receive the episodic returns $(G_{k,i+1}^{(1)}, G_{k,i+1}^{(2)})$.
 - 9: Compute reflection response rating by $r(x_{k,i}, y_{k,i}) \triangleq G_{k,i+1} - G_{k,i}$ in Eq. (5).
 - 10: Label the response with higher ratings as the accepted response while the lower response is labeled as the rejected response.
 - 11: **end for**
 - 12: **end for**
 - 13: **end for**
 - 14: **Step 2. Reward Model Learning.** Use the REWARDTRAINER in TRL to train a model for classifying accepted and rejected responses given instructions.
 - 15: **Step 3: Policy Gradient Finetuning.** Plug-in the trained reward model and use the PPOTRAINER in TRL to finetune the Retrospective model for generating reflection responses with higher ratings.
-

Fig. 12: Diagram outlining the RLHF-based training process, consisting of three steps: Obtaining training data; Training a reward model with Supervised Learning; and fine-tuning the *retrospective model* through response ranking using PPO [extracted from [13]].

B.2 Architecture Description using FALAA

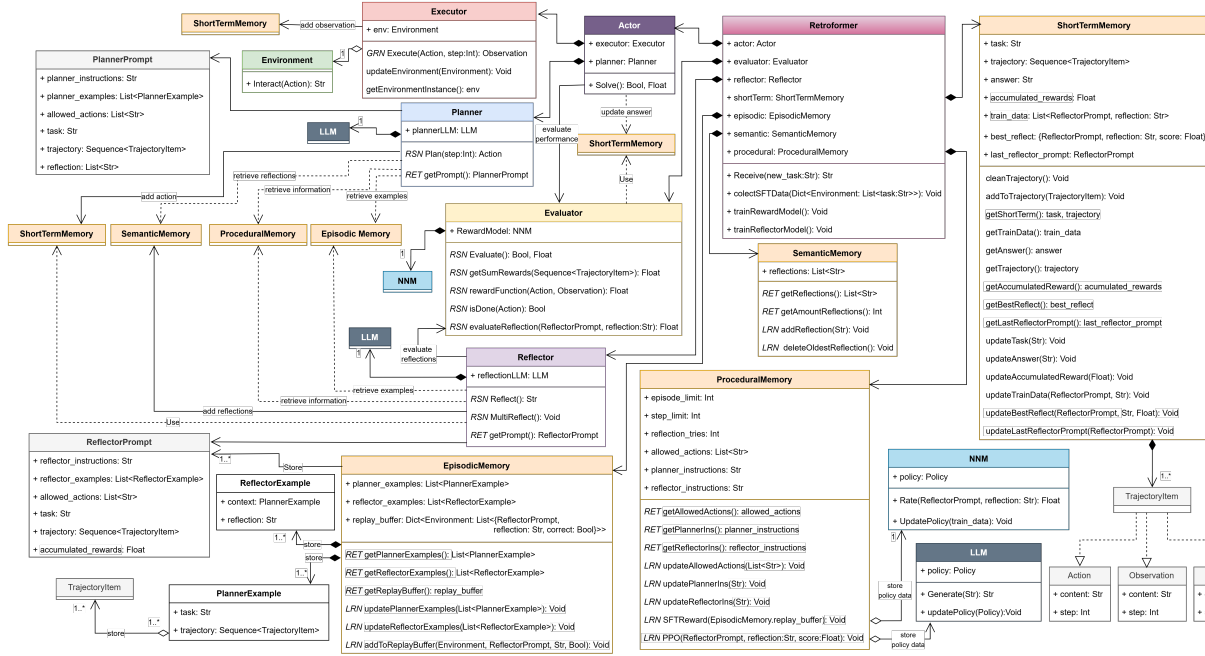


Fig. 13: UML class diagram of a Retroformer agent.

(1) Conceptual Description Level Components Description. Retroformer [13] originally proposes three fundamental components: an *Actor*, a *Retrospective model* and a memory module (short-term and long-term memory, as well as a *replay buffer*). Along with this components, a reward model is trained and used. Figure 13 shows the UML diagram exemplifying how these components are integrated under the standard structure proposed by FALAA. Each component is presented below, providing the original definition and its respective standardized adaptation:

1. Actor

- *Original Definition:* The *Actor* is described, like in *Reflexion*, as the component responsible for solving the assigned task, using an LLM to generate and execute actions in an environment.
- *Adaptation in FALAA:*
 - Under the standardized structure, it is defined as an additional component, consisting of two main subcomponents: the *Planner*, which is responsible for planning actions, and the *Executor*, which executes them.
 - The original concept of “solving the task” is maintained, but adapted to fit within the FALAA structure for better integration and compatibility with other agents.

2. Retrospective model

- *Original Definition:* It consists of a smaller LLM compared to the *Actor* and is responsible for producing reflections (*feedback*) when the *Actor* fails in the task.
- *Adaptation in FALAA:*
 - It is associated with the *Reflector* component, whose function is to generate reflections and insights on failures or possible improvements in the *Actor*’s strategy.

3. Reward Model

- *Original Definition:* Specifically trained by *Retroformer* to evaluate the quality of reflections generated by the *Retrospective model*.
- *Adaptation in FALAA:*
 - It is designated as an independent object within the architecture, similar to LLMs, dedicated to evaluation and reward assignment.

4. Memory Module

– *Original Definition:*

- *Short-term memory:* Stores the trajectory (actions, observations, rewards) generated by the *Actor*.
- *Long-term memory:* Saves reflections produced by the *Retrospective model*.
- *Replay buffer:* Specialized memory for storing triplets of **prompts**, reflections, and the accumulated reward of the episode.

– *Adaptation in FALAA:*

- **Short-term memory:** Under the standard structure, it stores a series of attributes, including the trajectory generated by the **Actor** and the accumulated rewards of that trajectory, as well as attributes aimed at the training process of **Retroformer**.
- **Long-term memory for reflections:** Reflections are stored in the **semantic memory** due to its nature of storing knowledge.
- **Replay buffer:** This particular memory is established as a dictionary that stores a triplet of prompts, reflections, and their classification as correct or incorrect for various environments. It is stored in the **episodic memory** due to its nature of storing past experiences.

Retroformer utilizes all structures proposed by FALAA, with the **procedural memory** playing a key role due to the architecture’s training processes. This memory in the **Retroformer** architecture it can execute *learning actions* through methods that implicitly update stored information, modifying the LLM and neural network-based reward model policies.

The training process also influences the **short-term memory**, which stores various information such as *train_data*. The trajectory stored in the **short-term memory** also undergoes modifications, adopting a sequence that now includes the elements **Action**, **Observation**, and **Reward**, reflecting the new architectural needs for a more structured and efficient process.

We define several Objects representing prompts and examples used in the creation of actions and reflections by **Planner** and **Reflector** respectively. This is to make their composition explicit, as in implementations they are often text strings or token lists.

The **Evaluator** component plays a crucial role within the architecture. Since it is responsible for generating evaluations and thus rewards, it integrates the reward model dedicated to evaluating the quality of reflections. **Evaluator** offers various functional methods, including the evaluation of an action-observation pair through a reward function, calculating the sum of rewards in a given trajectory, determining whether an action corresponds to a final response, and assessing the quality of a reflection.

In general, the architecture of **Retroformer** partially retains the structure of **Reflexion**, but introduces modifications and new functionalities, such as the use of rewards for each step of the trajectory or new attributes in the **short-term memory** used in the training processes of the reward model and retrospective model.

Behavior Description. Figure 14 shows the main UML sequence diagram of a **Retroformer** agent. Before interacting with the user, **Retroformer** exists in a certain environment, so its memories must be initialized (this fragment is identical to **Reflexion** in Figure 4). Once initialized, an infinite loop begins where the agent receives a new task from the user, which is added to the **short-term memory** before proceeding with the work cycle (condensed under the *ref* fragment referencing Figure 15). After completing this process, the generated response is retrieved from the **short-term memory** and returned to the user. Similar to **Reflexion**, it is not explicitly described what should happen if the response is not found, so we assume that the agent does not perform any additional

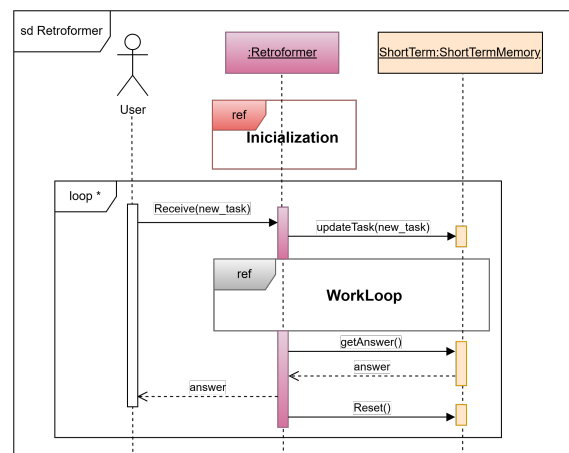


Fig. 14: UML main sequence diagram of the Retroformer LLM-based agent.

actions. Finally, the `short-term` memory is cleared using its `Reset()` method, leaving the agent ready to receive a new task.

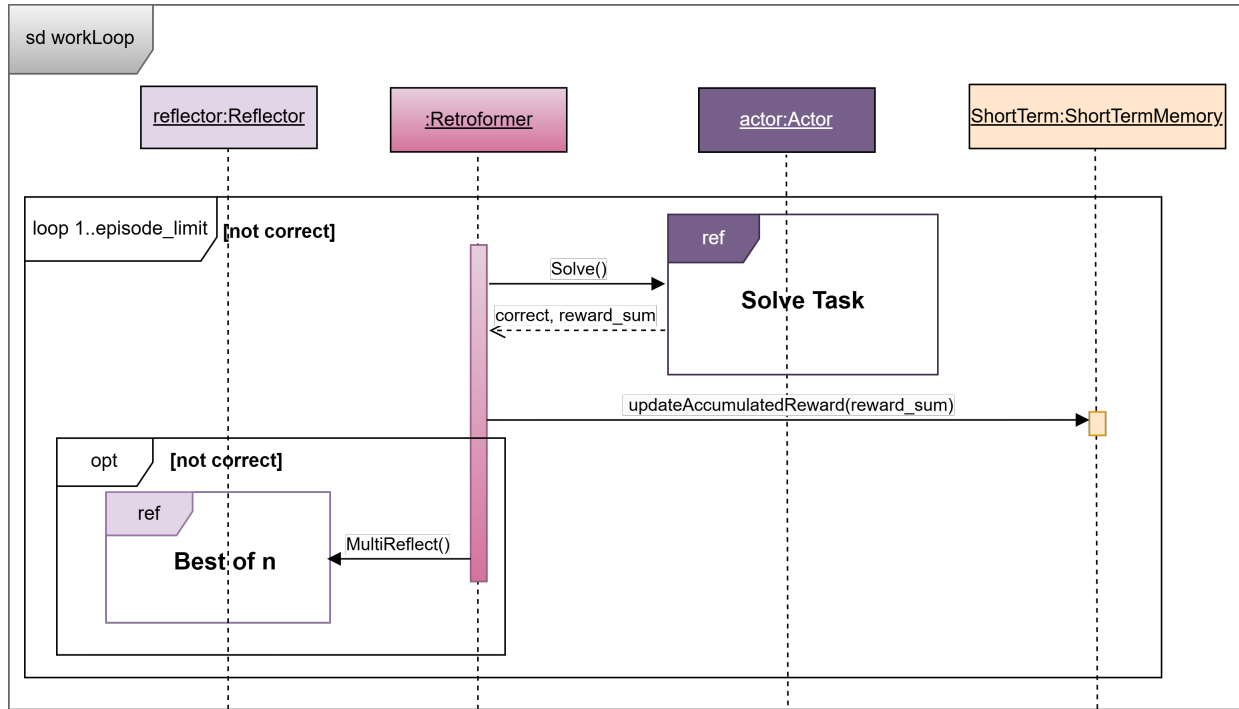


Fig. 15: UML sequence diagram of the work cycle of a `Retroformer` agent.

Regarding the referenced work cycle (see Figure 15), it is observed that the `Retroformer` agent maintains a work cycle similar to that of `Reflexion`, consisting of a loop running from 1 to the maximum value `episode_limit`, which indicates the number of failed responses allowed before ending the cycle. The loop continues as long as the response generated in the previous iteration is not correct. In each iteration, the `Actor` attempts to solve the task (see Figure 16), returning the sum of the accumulated rewards in its trajectory along with an indication of whether the task was successfully resolved. If the task is not successfully resolved, the process of generating reflections is executed, introducing a new flow (see Figure 18). Finally, a new iteration begins.

The task resolution process of the `Actor`, as shown in Figure 16, remains as general as possible, since `Retroformer` allows replacing the `Actor` with components from other agents. The task resolution flow is based on two parts.

The first part is a loop with a maximum of `step_limit` steps, where in each step, `Planner` is asked to generate an action (the planning and corresponding information retrieval process for prompt creation is the same as in `Reflexion`, see Figures 7 and 10, respectively). The generated action is then evaluated to determine if it is the final response, which would end the loop. If not, `Executor` executes the action, receiving an observation that is added to the trajectory in the `short-term` memory (the process is the same as in `Reflexion`, see Figure 8). Then, the action-observation pair is evaluated using the reward function of `Evaluator`, returning a score for this step, which is converted into a `Reward` object and added to the `short-term` memory.

The second part occurs when the loop ends, either by reaching the step limit or by finding the correct response. The current episode is then evaluated, where the accumulated rewards from the current trajectory are obtained, and it is determined whether the response is correct or not, returning these two values.

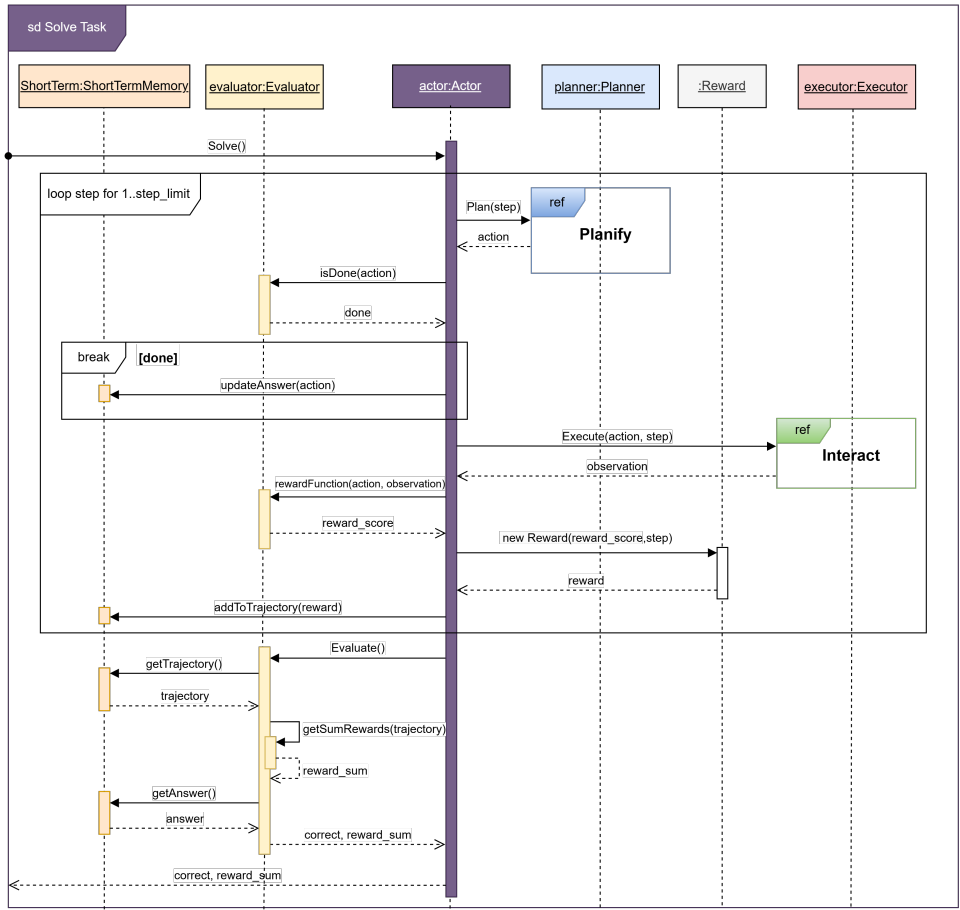


Fig. 16: UML sequence diagram of the task resolution process by the Actor component of a Retroformer agent.

Figure 18 shows the process of selecting the best reflection from the Reflector component of a Retroformer agent, which is the main innovation of this architecture. It illustrates how Reflector generates *reflection_tries* reflections in a loop. In each step of this loop, a reflection is generated using the `reflect()` method (see Figure 17), where a prompt composed of information from short-term and long-term memories (see Figure 19) is used to generate a reflection, which is then returned.

Once the reflection is generated, it is evaluated using the Evaluator component, which acts as an intermediary. Using the prompt and the generated reflection, Evaluator applies its pre-trained neural network model to assess the quality of the reflection. If the reflection is better than the best one stored so far, it is saved as the best reflection in the short-term memory. At the end of the cycle, the best reflection is stored in the semantic memory, ensuring that the storage limit is not exceeded.

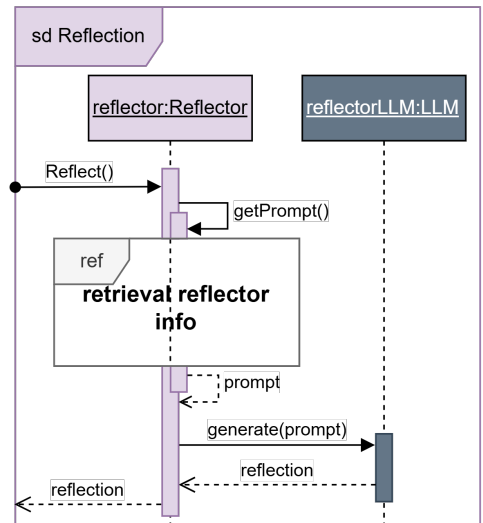


Fig. 17: UML sequence diagram of the reflection process of the Reflector component of a Retroformer agent.

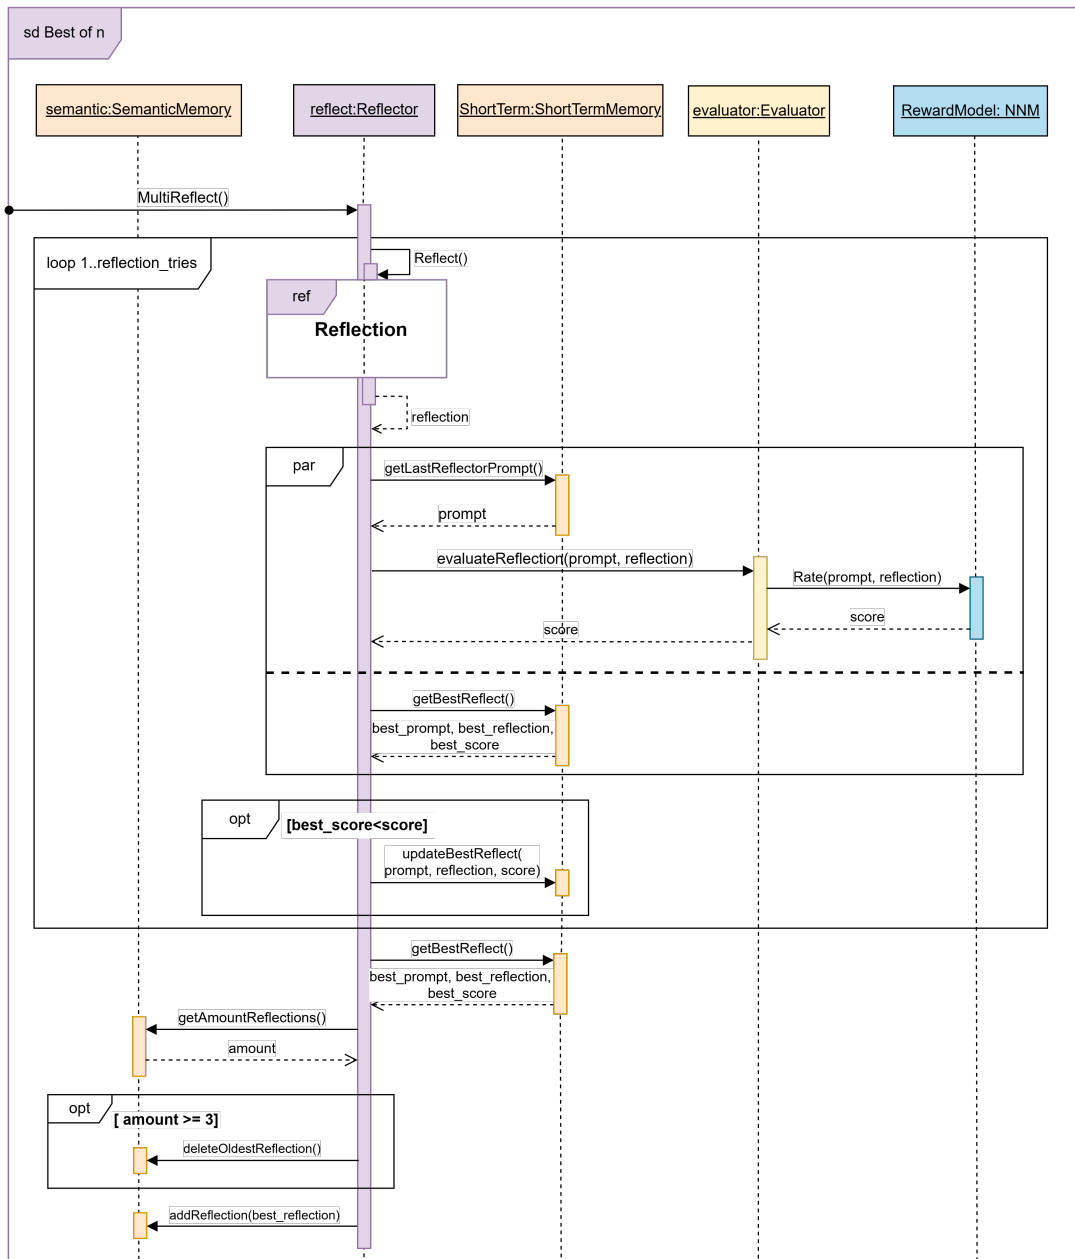


Fig. 18: UML sequence diagram of the best reflection selection process for the Reflector component of a Retroformer agent.

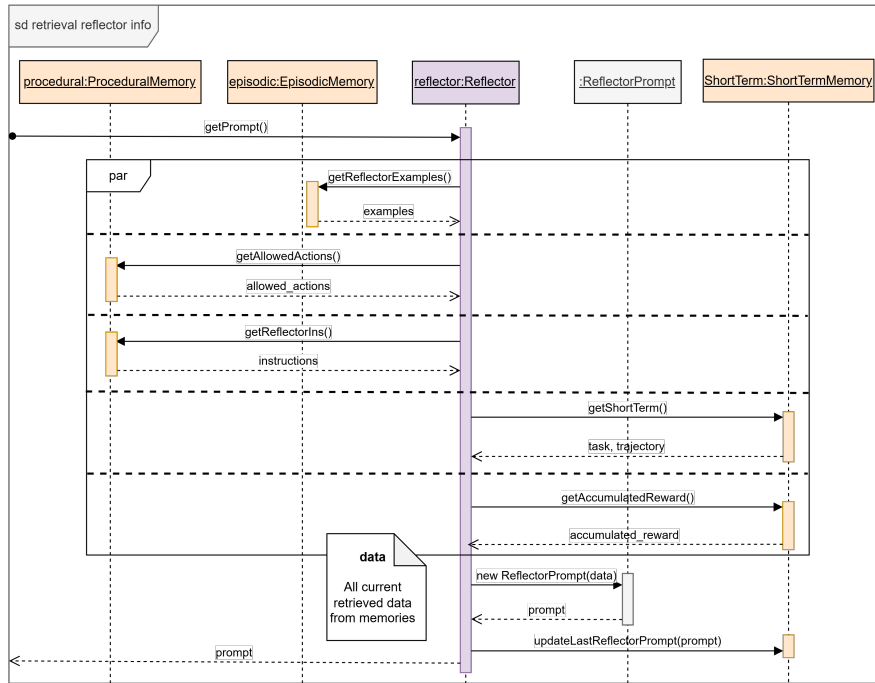


Fig. 19: UML sequence diagram of the information retrieval process for reflection in the **Retroformer** agent. In this flow, **Reflector** uses internal **retrieval** actions, represented by the *get* methods of the **episodic** and **procedural** memories, from which examples for the LLM and the allowed actions along with instructions to guide the LLM’s behavior are obtained. These, together with the current task and trajectory retrieved from the **short-term** memory, are used to generate the prompt. This prompt, before being returned, is stored in the **short-term** memory.

Training Process. **Retroformer** employs an RLHF-based training method [5], where it trains a neural network model using supervised learning to evaluate reflections, along with the supervised learning-based training of the LLM contained in **Reflector**. Figure 20 illustrates the general training process of a **Retroformer** agent. This process consists of the three RLHF steps, which are abstracted into the methods *collectSFTData(train_data)*⁶, *trainRewardModel()*, and *trainReflectorModel()* of the **Retroformer** agent. The three sequences are referenced in the diagram and correspond to Figures 21, 24, and 25, respectively.

The first step of the RLHF-based training focuses on data collection. Figure 21 shows how, after receiving a dictionary with environments and their respective tasks, the **Retroformer** agent iterates over each environment, updating the environment instance stored in the **Executor** component. Then, memories are initialized based on this new environment (the same process as in **Reflexion**, see Figure 4). Each task associated with the environment is iterated a total of three times. For each task, the process is quite similar to the work cycle of a **Retroformer** agent (see

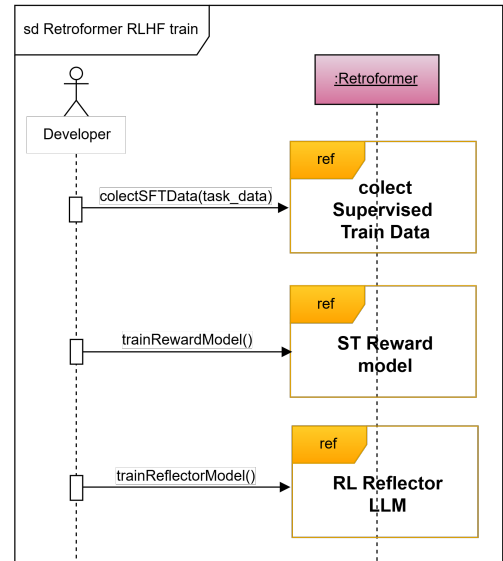


Fig. 20: Main UML sequence diagram of the RLHF-based training process of a **Retroformer** agent.

⁶ The information contained in the *task_data* variable is based on a dictionary where the key represents various environments, and the value associated with each key is a set of tasks. A detailed view of the types can be found in the component description in the class diagram in Figure 13

Figure 15), where the task is solved, the response is evaluated, and the information is stored in the **short-term** memory. The difference is that if the response is incorrect, a process is initiated to generate two reflections (see Figure 22), which are then evaluated by the **Evaluator** and stored in the *replay buffer* (see Figure 23) of the episodic memory. Finally, the **short-term** memory is reset, and the next task is processed.

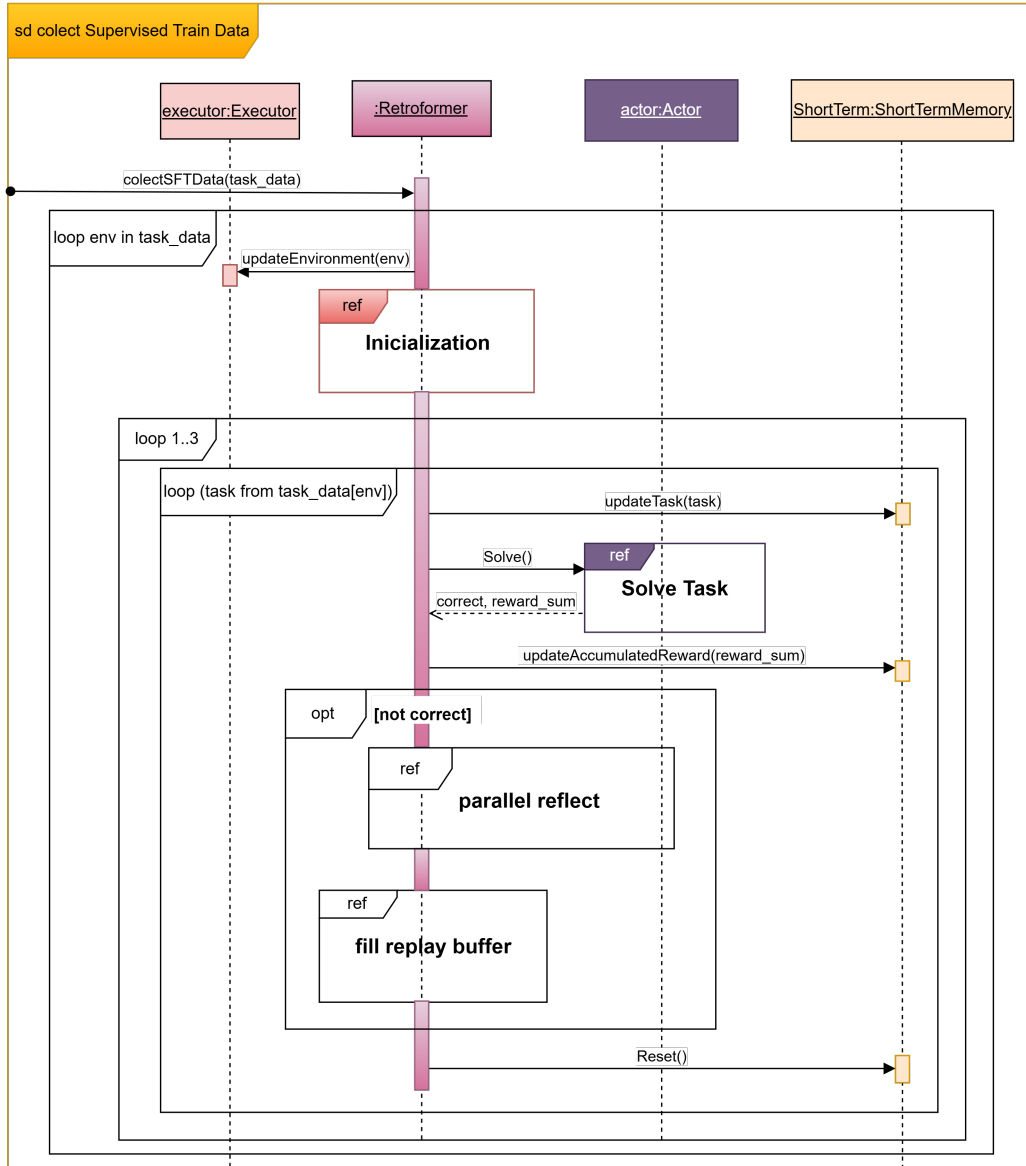


Fig. 21: UML sequence diagram of Step 1: Data collection for a **Retroformer** agent.

The process depicted in Figure 22 is another proposal from [13]. To avoid human intervention in labeling good and bad reflections, the authors of **Retroformer** propose using the difference between the accumulated rewards of two successive episodes as a signal to classify reflections into good and bad categories.⁷

⁷ This is because the **Planner** LLM is considered a frozen parameter model with a low temperature. In this context, temperature refers to the parameter that configures the randomness in the LLM's response generation, implying that randomness injection is minimal. Therefore, changes in accumulated rewards are mainly attributed to the **Reflector** LLM.

In this process, based on the same trajectory, two reflections are generated. Then, each reflection is iterated, and the task is attempted again, obtaining accumulated rewards for each episode. Finally, for each of the two reflections, the prompt, reflection, and new accumulated rewards are stored in the **short-term memory** for further classification.

The final part of Step 1: Data collection is illustrated in Figure 23, which shows how the generated reflection information is stored in the *replay buffer* of the **episodic memory**. The addition of this data is based on the difference in accumulated rewards between the two consecutive episodes obtained for each reflection. Comparing this score determines which reflection is correct and which is incorrect. An undefined case by [13] is highlighted, where both scores obtained for the reflections are equal. In such a case, both reflections are classified as correct; however, it is noted that this remains an open issue.

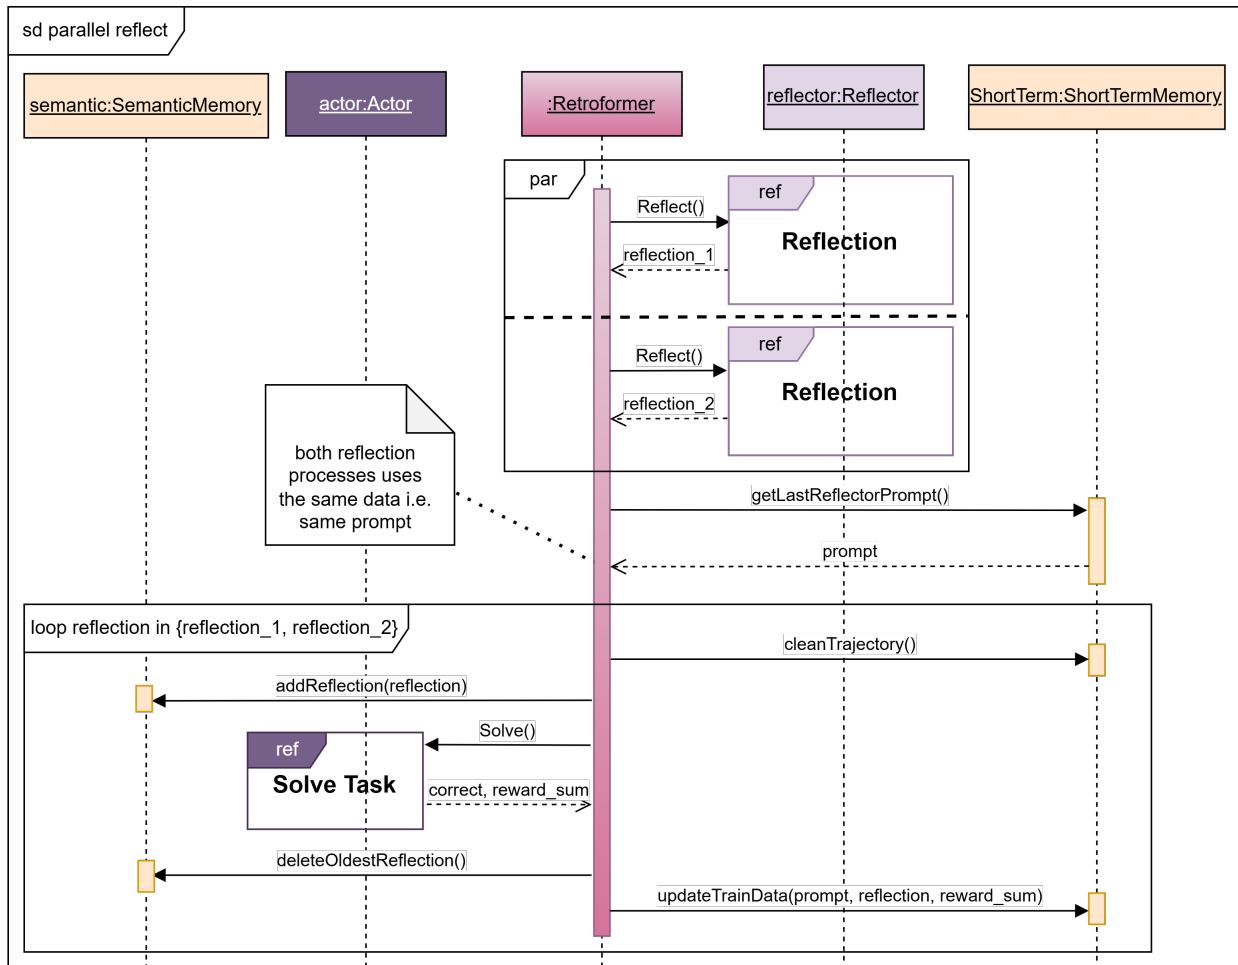


Fig. 22: UML sequence diagram of the flow that creates training data based on two reflections of a **Retroformer** agent.

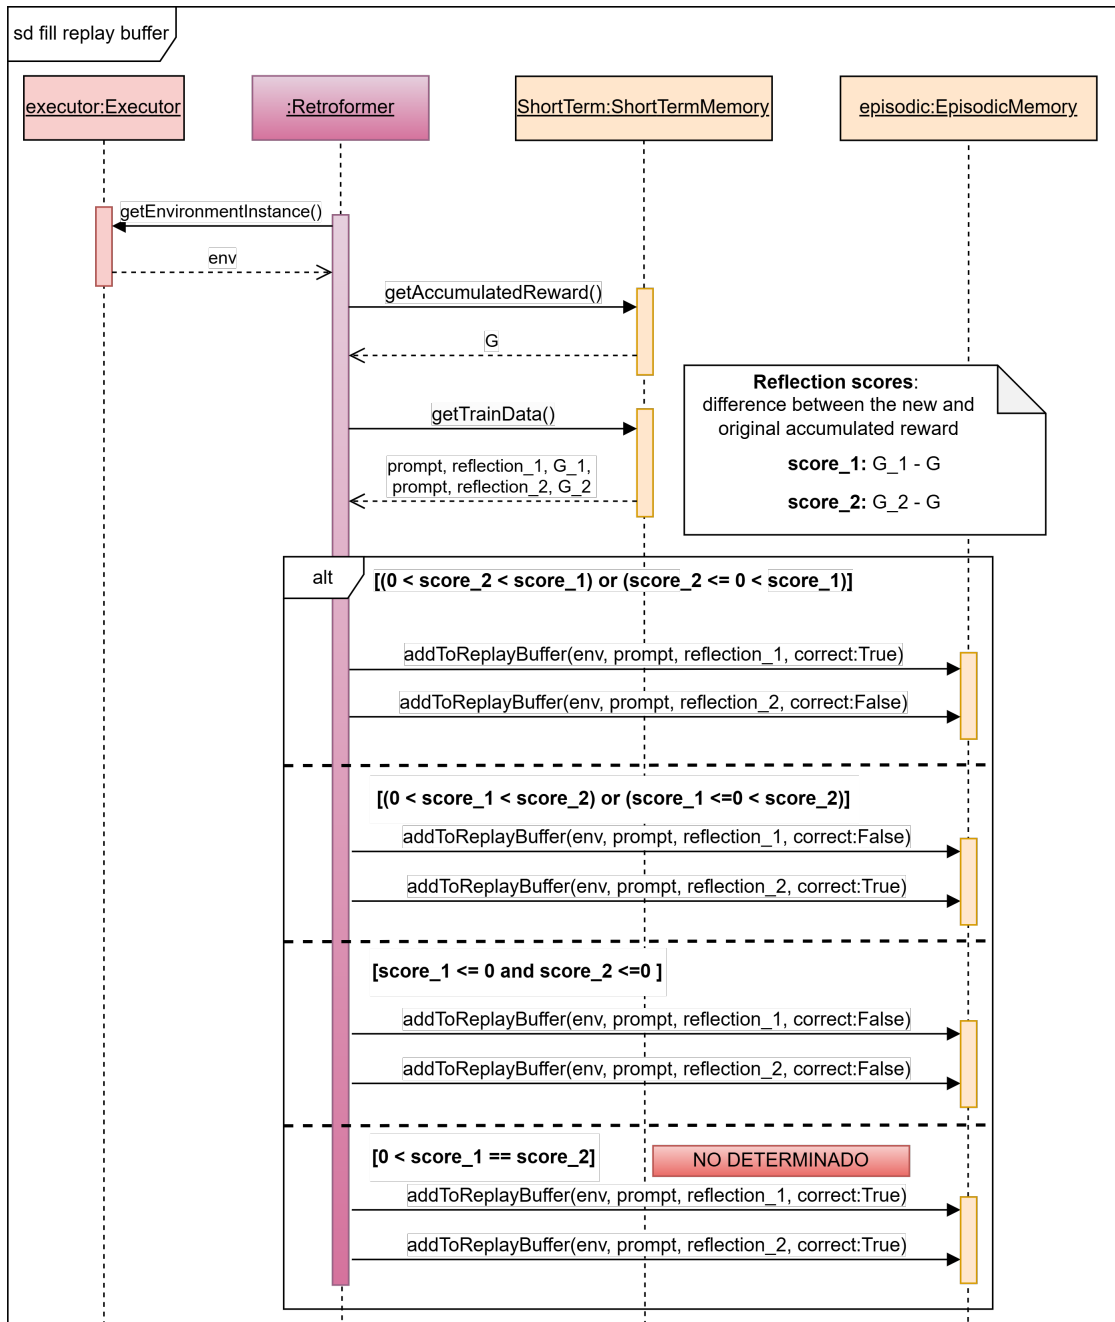


Fig. 23: UML sequence diagram showing the process of filling the *replay buffer* of the episodic memory in a Retroformer agent.

Step 2 of the training process for a Retroformer agent, illustrated in Figure 24, focuses on training a reward model using information from the *replay buffer* of the episodic memory. In this process, Retroformer utilizes the procedural memory (which implicitly stores the information contained in the parameters of both the reward model and the LLMs) to update the reward model’s policy. The reward model is trained using supervised learning based on the data stored in the *replay buffer*.

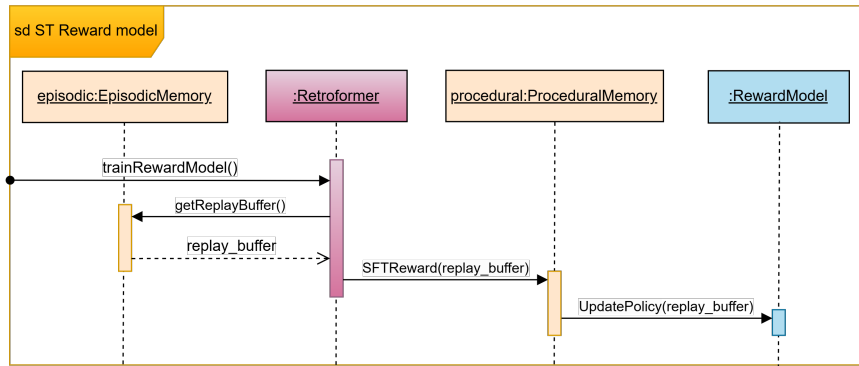


Fig. 24: UML sequence diagram of Step 2 in RLHF-based training of a neural network reward model for a Retroformer agent.

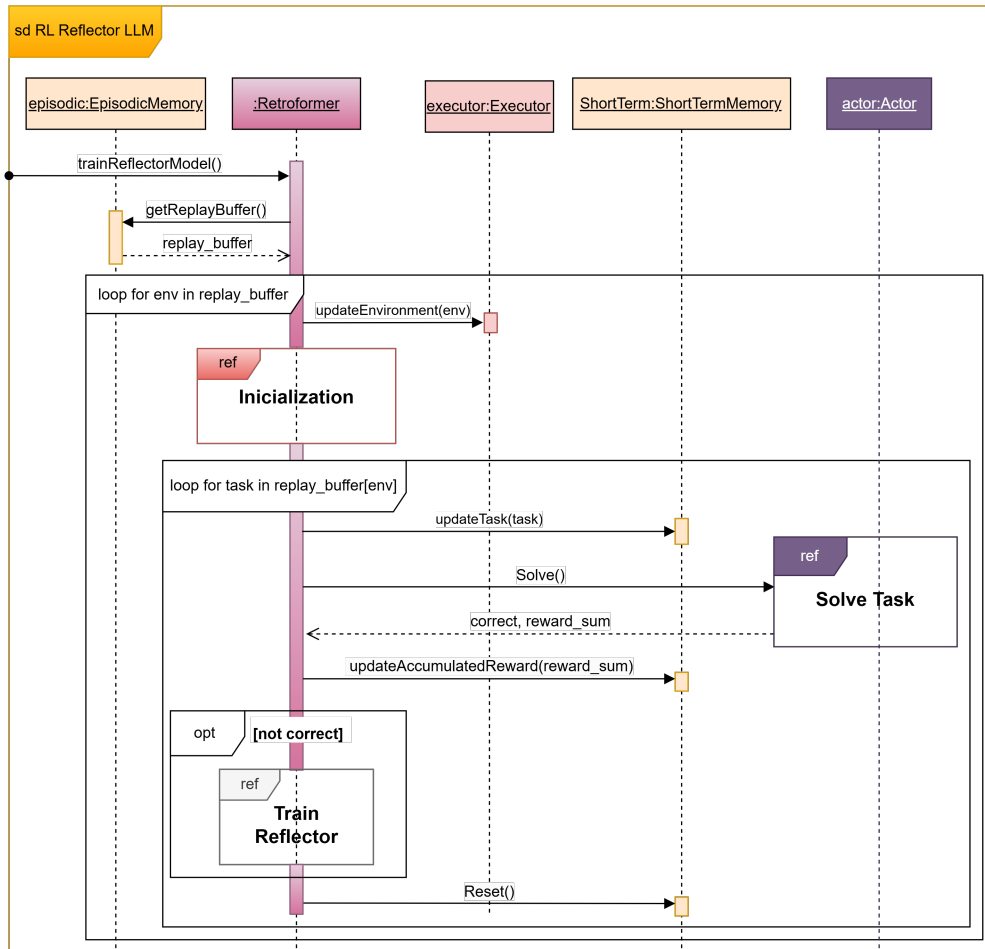


Fig. 25: UML sequence diagram of Step 3 in RLHF-based training of a Retroformer agent, where the LLM of the Reflector is trained using reinforcement learning.

Finally, Step 3 of the RLHF training process is depicted in Figure 25. In this step, the LLM of the Reflector is trained using reinforcement learning, leveraging the reward model trained in the previous step. The process iterates over each environment and its respective tasks stored in the *replay buffer* of the episodic memory.

Similar to Step 1 of data collection, the environment is initialized, and the task is solved using the Actor. The difference is that now, when a response fails, the process shown in Figure 26 is triggered, where a reflection is generated, evaluated, and the LLM of the Reflector is fine-tuned based on this reflection using the PPO algorithm [6].

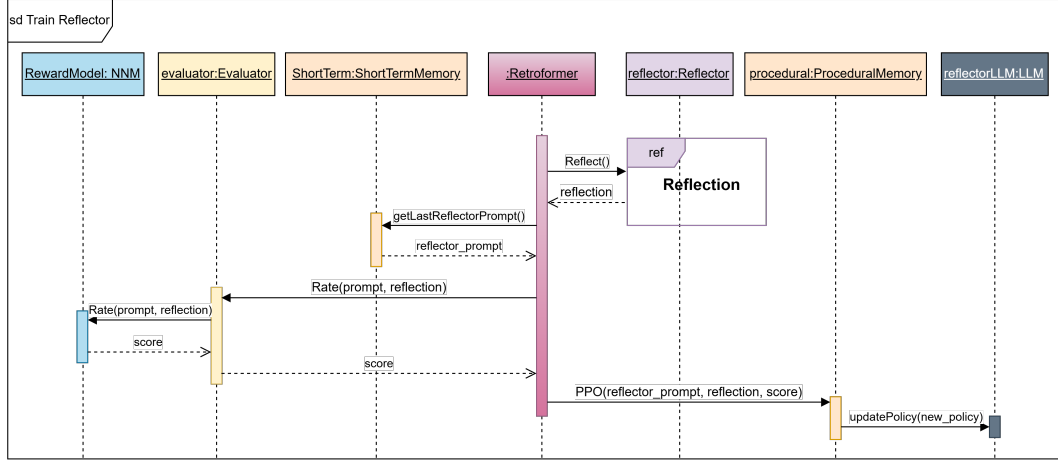


Fig. 26: UML sequence diagram where, given a reflection, it is evaluated using the reward model of Evaluator, and this score is then used to update the policy of the LLM in the Reflector component using the PPO algorithm.

(2) **Formal Specification Level** In the Retroformer architecture, we have identified seven behaviors that, to avoid ambiguities, need to be formally specified. These behaviors are:

1. The *Reset* method of the `short-term memory`, which must clear its attributes in a specific format.
2. The *isDone(Action)* method of `Evaluator`, which must return true if the received action contains the string “[answer]”.
3. The `semantic memory` can store at most 3 reflections.
4. The respective precondition of the *addReflection* method in `semantic memory`, which allows adding a new reflection only if there are fewer than 3 reflections stored.
5. The trajectory invariant of the `short-term memory`, which states that stored objects must be added in the sequence Action-Observation-Reward.
6. The pre and postconditions that the *addToTrajectory* method of `short-term memory` must follow.
7. In the training process of `Retroformer`, 2 reflections are used to generate training data. We decided to explicitly define this invariant, where the information temporarily stored in the *train_data* attribute of the *short-term memory* consists of 2 reflections, one correct and one incorrect.

The behaviors 1, 2, 3, and 4, specified using OCL, are described similarly to their counterparts in `Reflexion`, which are presented in equations 1, 2, 3, and 4, respectively.

The specifications of behaviors 5, 6, and 7 are presented below in equations 7, 8, and 9, respectively.

```

context ShortTermMemory inv :
  self.trajectory → notEmpty()
  or
  let size: Integer = self.trajectory → size() in
  Sequence{1..size} → forAll(i :Integer |
    if i.mod(3) = 1 then
      self.trajectory → at(i).oclIsTypeOf(Action)
    else if i.mod(3) = 2 then
      self.trajectory → at(i).oclIsTypeOf(Observation)
    else
      self.trajectory → at(i).oclIsTypeOf(Reward)
    endif
  )

```

(7)

```

context ShortTermMemory :: addToTrajectory(newItem : TrajectoryItem): OclVoid
pre :
  if self.trajectory → isEmpty() then
    newItem.oclIsTypeOf(Action)
  else let lastElement = self.trajectory → last() in
    if lastElement.oclIsTypeOf(Action) then
      newItem.oclIsTypeOf(Observation)
    else if lastElement.oclIsTypeOf(Observation) then
      newItem.oclIsTypeOf(Reward)
    else if lastElement.oclIsTypeOf(Reward) then
      newItem.oclIsTypeOf(Action)
    endif
  post :
    self.trajectory → size() = self.trajectory@pre → size() + 1
  post :
    self.trajectory → last() = newItem

```

(8)

```
context ShortTermMemory::Reset(): OclVoid
post:
  self.task = ""
post:
  self.answer = ""
post:
  self.trajectory → isEmpty()
post:
  self.accumulated_rewards = 0
post:
  self.train_data → isEmpty()
post:
  self.best_reflect → isEmpty()
post:
  self.last_reflector_prompt → isEmpty()
```

```
context ShortTermMemory inv :
  self.train_data → size() <= 2
```