



HAL
open science

Leveraging Large Language Models for enhanced personalised user experience in Smart Homes

Jordan Rey-Jouanchicot, André Bottaro, Eric Campo, Jean-Léon Bouraoui,
Nadine Vigouroux, Frédéric Vella

► **To cite this version:**

Jordan Rey-Jouanchicot, André Bottaro, Eric Campo, Jean-Léon Bouraoui, Nadine Vigouroux, et al..
Leveraging Large Language Models for enhanced personalised user experience in Smart Homes. 2024.
hal-04593247

HAL Id: hal-04593247

<https://hal.science/hal-04593247>

Preprint submitted on 27 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Leveraging Large Language Models for enhanced personalised user experience in Smart Homes

Jordan Rey-Jouanchicot^{1,2,3}, André Bottaro³, Eric Campo², Jean-Léon Bouraoui³, Nadine Vigouroux¹, Frédéric Vella¹

¹IRIT, University of Toulouse, CNRS, Toulouse INP, UT3, UT2J, Toulouse, France

²LAAS-CNRS, University of Toulouse, CNRS, UT2J, Toulouse, France

³Orange Innovation, Blagnac, 31700, France

Abstract

Smart home automation systems aim to improve the comfort and convenience of users in their living environment. However, adapting automation to user needs remains a challenge. Indeed, many systems still rely on hand-crafted routines for each smart object.

This paper presents an original smart home architecture leveraging Large Language Models (LLMs) and user preferences to push the boundaries of personalisation and intuitiveness in the home environment. This article explores a human-centred approach that uses the general knowledge provided by LLMs to learn and facilitate interactions with the environment.

The advantages of the proposed model are demonstrated on a set of scenarios, as well as a comparative analysis with various LLM implementations. Some metrics are assessed to determine the system's ability to maintain comfort, safety, and user preferences. The paper details the approach to real-world implementation and evaluation.

The proposed approach of using preferences shows up to 52.3% increase in average grade, and with an average processing time reduced by 35.6% on Starling 7B Alpha LLM. In addition, performance is 26.4% better than the results of the larger models without preferences, with processing time almost 20 times faster.

1 Introduction

Networks of devices are deployed to assist human beings in their daily activities, using notions of context and knowledge to decide on the best actions to take [Weiser, 1991] [Dey, 2001]. Indeed, many houses are covered by wireless and wired networks and equipped with electronic devices allowing the occupants to control their environment for comfort, entertainment, security, energy management, and elderly care.

However, Smart Home Automation Systems are still missing the aim of autonomously taking the best action in every situation. Aligning automation routines to meet every need

for every home configuration, every set of devices, available or not, functional or not, remains a challenge. In fact, most systems today are configured for simple routines that occur frequently. For instance, setting the home for wake-up or departure time by playing music, ringing a bell, acting on lights, shutters, heating, ventilation, and air conditioning. More precise needs in less frequent situations are not covered. Furthermore, if some devices are missing from a routine, no decision is made to use alternative devices. While artificial intelligence can play a role in learning about situations and the associated actions taken by users [Rashidi and Cook, 2009], it still requires time to understand users' habits, and is never able to cover the wide range of situations encountered at home.

One of the main technical challenges of pervasive computing is the ability to set up a system knowing the wide variety of users' potential needs, and how it can adapt to the wide range of functions of existing devices, with devices and locations whose configuration can be very different. As their execution environment is particularly dynamic, applications need to be aware of their context and act appropriately.

Large language models [Radford *et al.*, 2019] have the potential to give this general knowledge at once to applications such as smart home. In essence, these models have acquired a vast amount of knowledge. They are trained on a diverse range of textual sources, enabling them to cover a variety of topics, facts, and concepts, here the expected actions of home devices to meet user needs in a wide range of situations.

Retrieval-Augmented Generation [Lewis *et al.*, 2020] is a technique that improves the accuracy of LLMs by giving them access to more targeted and precise information. This is achieved by using a retrieval component that searches a specific database and introduces it into the model, along with users' query.

This paper proposes a software architecture integrating the general LLM knowledge available today into a smart home automation system. The LLM is placed at the center of the home's decision-making system and participates in the reaction to every event to deduce the next best actions. This paper investigates the inclusion of preferences with a LLM for smart home automation, The contribution also includes a user-centred representation of smart home states and actions in natural language. Finally, experiments are carried out using several LLMs with different prompting styles for decision-making in the smart home.

The next section presents related work in the literature. Then, the proposed software architecture is detailed, in particular the integration of LLMs for decision-making combined with RAG for user context retrieval. Afterwards, the experimentation section describes a dedicated benchmark of LLMs and prompting styles with a set of home scenarios, leading to the results in the next section. These results are discussed and a final section is devoted to the conclusion and future work.

2 Related works

This topic is recent and few papers have been published on LLMs for smart homes. The following papers are selected for their approaches using the knowledge of user preferences on decision-making in smart home automation systems. This section cannot be exhaustive in this larger domain.

[Oliveira *et al.*, 2022] proposes a multi-agent environment with a Belief-Desire-Intention cognitive model, to support adaptivity and preferences transparently in a smart home environment. Any possible interaction could be modeled in such a way as to allow alternative proposals, but this requires considerable work in semantic representation to be fully complete.

Another paper takes advantage of general knowledge information to improve system adaptivity to preferences. [Ruiz *et al.*, 2021] proposes an approach using three Knowledge-Based systems, one with general knowledge, one with skill knowledge, and one with contextual information such as device location, and generating rule models for a middleware platform based on all this information.

This approach requires really strict semantic modeling to handle most scenarios, and cannot take advantage of some information that is implicitly given in context.

[Shuvo, 2022] proposes an actor-to-critic (A2C)-based algorithm adapted to decision-making in smart homes for energy consumption. In this work, at each step, an A2C algorithm is applied to each device to select the best action, using as inputs the activity and price of electricity at that time.

A key element is that the set of actions for each appliance depends on the category associated with the appliance, adding initial knowledge to the model to ensure action based on the importance of the appliance. This system requires learning, and any change in device availability leading to different optimal decisions will require a large number of steps before adapting.

In addition, the larger the context, the more difficult it will be to converge for each scenario to learn the best-suited device state. The paper HESIDR, [Zhang *et al.*, 2023] also proposes a system for energy consumption in smart homes, it combines a set of control rules and reinforcement learning to reduce the adaptation time.

The study proposed by [Peng *et al.*, 2020] describes an approach for decision-making in home automation using deep reinforcement learning. It showed the ability to learn when to turn on a light but with a really limited context supported, so the application is limited to learning when to turn on a light with schedules of 15 minutes, which is a quite large period for this application.

[King *et al.*, 2023a; King *et al.*, 2023b] proposes the first

approach to smart home automation, using a JSON data representation from a smart home middleware platform and experimenting with an LLM to select an action based on a user request.

This approach is a first step towards the use of the general knowledge provided by these models. However, it does not support user preferences, and the idea is to select actions based on an initial user request, rather than proactively. These works have proposed ways of managing decision-making, but are limited by the contextual data supported. Works using symbolic AI methods can show great adaptability, but at the cost of extensive semantic modeling and with requirements to make them adaptable to future changes in preferences. A new method is proposed to support contextual data while adding preferences.

3 Proposed architecture

This study proposes a new architecture for decision-making in smart home automation systems. The system uses Large Language Models and proposes methods to add user preferences, in order to select an action according to context and users. It aims to be a proactive system. At every event occurring in the home, the system proposes actions on devices to align the home state with user needs and preferences. The system supports different types of data thanks to LLMs ability to process data while generating a textual representation of the home based on device states and the action list.

The proposed system aims to be able to adapt to changes in user preferences over time. Indeed, LLMs avoid the need for retraining and handle appliance configuration changes, basic unavailability, or appliance failure.

This reason leads directly to the use of RAG or directly injecting the knowledge into the prompts for retrieving updated preferences at every execution time of the AI, instead of fine-tuning. Indeed, fine-tuning over preferences would be impractical, as it would require a new iteration each time user preferences change, which would be costly and computationally expensive.

A home simulator is implemented, it takes information on the configurations of sensors and then generates a textual representation, it is also used to generate the list of actions from its data.

This section details the main components designed for the proposed system architecture. Figure 1 shows this architecture, including the simulator.




Figure 1: The implemented architecture

The system generates a user-centred text description of the

home and a list of actions, with control over connected devices. It filters the number of relevant actions that can be taken in every situation. For example, it limits some actions that may be prohibited to guarantee user safety. This list of actions is used by the model to select the optimal action.

Concerning contextual data representation editing:

- User positions are listed with their current activity and the history of previous activities.
- The history of previous actions performed in the house is supported.
- All rooms, sensors, and actuators are presented using their names, which the user optionally gives.
- Some sensors and actuators give more global data and control, e.g., internal and external temperature sensors, gas sensors, humidity sensors, HVAC systems, etc.

The implementation supports some device categories with a dedicated natural representation, to generate more natural sentences adapted to some types of data: lights, CO2 sensors, smart curtains, etc. It could also support any additional data sensors with a generic representation template, using the device name in the smart home environment and data status. Using meaningful naming added by the users helps the system to understand the usage of the device.

The action proposal algorithm 1 considerably reduces the set of possible actions. It assumes that only devices that are in the room or global can be switched on, but that all switched-on devices can be switched off. As far as the list of possible actions is concerned, the idea is to filter the actions supported according to some conditions and device types. This approach is made possible by LLM’s native support for a change in the output action space, without requiring training.

Algorithm 1 Dynamic Devices: Action builder Algorithm

Require: *userid, devices_list*

```

1: for all device_name, device_kind, device_location,
   device_state in devices_list do
2:   if device_kind = "actuator" then
3:     if user_location[userid] = device_location or
       device_state = 1 then
4:       devices.append(device_name is device_state)
5:       action_vector.append(1)
6:     end if
7:   end if
8: end for
9: devices.append("Interact with user")
10: action_vector.append(2)
11: devices.append("No action required")
12: action_vector.append(0)
13: return (action_vector, devices)

```

Any type of device can be easily supported: it simply has to be added to the representation and the LLM will ingest the data thanks to its internal knowledge. This knowledge allows the LLM to get a natural human description of the home including biased contextual data unlike many conventional home automation systems. The latter do not take advantages

of information such as the names of lights or rooms. This data is transmitted to the LLM using one of the prompting styles, prompting being the way to call the LLM with contextual arguments. The different styles of prompts will be described in the following section.

A common aim for all the prompting styles is to take advantage of the knowledge of the overall world provided by this Large Language model, to handle changes or even new types of sensors added to the representation.

The user preferences and rules block represents a database containing information about the system’s basic rules, generalities about human preferences and specific user preferences.

A benchmark is established with predefined scenarios for evaluation purposes.

Regarding LLMs, some off-the-shelf models are selected, with a local inference engine backend.

4 Experimentation

Different objectives are defined for the experimentations.

- Evaluate the improvements provided by adding user preferences, with various techniques of doing so.
- Evaluate the improvements of natural language representation of a smart home automation state over a JSON representation, as LLMs are trained on natural language corpus. Even if they contain other kinds of data like code.

In alignment with the objectives of experimentations, different metrics are used for these evaluations:

- Grades: The grade obtained for each model, the prompting styles for each scenario, by execution. , based on the grade values defined in table 1
- Processing time: Total runtime, including the construction of the context data and action list representation, the inferences with the prompting styles, the use of RAG if the prompting style uses it, and the processing of the formatted LLM response.

Eleven evaluation scenarios are defined as starting points, with predefined accepted actions and their specific grade. Table 1 presents the eleven scenarios by name and the associated reward values. A category is associated with each scenario, the goal being to regroup the scenario with the name of the main evaluated ability.

Figure 2 shows an example on scenario 1 of the generated contextual representation (textual) transmitted to the LLM.

The database of preferences and rules is defined in a single file for all scenarios. These data are naturally written sentences, and at the end of each one, information about the style is recorded: Rules, Preferences, Generality. The idea is to transmit to the LLM the importance of each data through keywords. Generality is considered the least important, Preferences the second most important, and Rules the most important. The database includes some preferences, generality and some rules. It is designed to handle some scenarios, help in some others but does not provide a solution for all scenarios. The data are fed into a vector database so that RAG may be used instead of prompts to convey them to the LLM.

Scenario Name	Grade	Associated Answer	Category
Out of bed at night	2	Turn on auxiliary light or main light with reduced luminosity level	Safety
	1	Turn on main light	
	0	Everything else	
Watching TV: late evening	2	Turn on auxiliary light or main light with reduced luminosity level	Comfort
	1	Turn on main light, open curtains, discuss	
	0	Everything else	
Out from bed issue with CO2	2	Inform user of risk	Safety
	1	Do an action and inform the user of risk	
	0	Everything else	
Going back to bed at night	2	Turn on auxiliary light or main light with reduced luminosity level	Safety
	1	Turn on main light	
	0	Everything else	
Evening sleeping: TV ON	2	Turn off TV	Preference
	1	Turn off anything on	
	0	Everything else	
At dinner watching TV	2	Turn on auxiliary light or main light with reduced luminosity level, open curtains	Preference
	1	Turn off the main light, do nothing	
	0	Everything else	
Forgot to turn off TV: user out	2	Turn off TV, turn off HVAC	Comfort
	1	Turn off all lights	
	0	Everything else	
Too low temperature	2	Turn on HVAC	Preference
	1	Open Curtains	
	0	Everything else	
Low luminosity day	2	Open curtains	Preference
	1	Turn on any light in the room	
	0	Everything else	
Failed curtains	2	Turn on any light of the room	Comfort
	1	Open curtains	
	0	Everything else	
Forgot to turn off lights	2	Turn off any lights, all lights, or HVAC	Preference
	1		
	0	Everything else	

Table 1: Scenario responses with grades, associated answers, and evaluation categories.

Four different prompting styles are compared on all the scenarios:

- **direct**: A system prompt and a prompt to request direct answers in the specified format.
- **directPref**: A system prompt with the preferences, rules and generality from the database and a prompt to request answers in the specified format.
- **OpenQuestion**: Two-steps chain: A system prompt, and a prompt to request: "a list of 3 main problems". For each of the 3 problems uses RAG to get the 3 closest preferences. A prompt to request answers in the specified format
- **ThreeQuestion**: Three-steps chain: A system prompt, a prompt to request: "a list of 3 main problems". For each of the 3 problems uses RAG to get the 3 closest preferences. A prompt to request answers in the specified format requests (2 times). A final prompt to request an answer in the specified format (based on the previous ones).

The database and prompting styles are available in the supplementary material.

A common point between all the prompting styles is the action expected in the output: as mentioned above, most supported devices, such as lights or HVAC systems, are considered as switches in the action list; therefore, in the output of

all prompting styles in addition to a "reasoning" and an "action" key, three optional keys are available: temperature, luminosity and explanation. It enables the model to respectively modify the temperature of an HVAC system when executing a related action, modify the luminosity by dimming a light or give an explanation to transmit a sentence to the user.

Two ways of representing the state of the house data are implemented, both using the same input data:

- **JSON**: A JSON representation
- **Textual**: A fully natural textual representation

The implementation of the system is evaluated using various open-sources LLMs, including:

- Starling Alpha 7B [Zhu *et al.*, 2023]- 8bpw
- Qwen 1.5 14B [Bai *et al.*, 2023] - 5bpw
- Qwen 1.5 72B [Bai *et al.*, 2023] - 3.5bpw

The three models are selected for their performance and for covering the three main open-source model sizes. They are used to evaluate the impact of proposed prompting methods and data representation.

Qwen 1.5 72B, with around 72 billion parameters, is currently one of the best models available open-source. Starling 7B Alpha, with around 7 billion parameters, is an excellent smaller model, and is based on Mistral 7B an efficient model for its size on various benchmarks. Qwen 1.5 14B model, a smaller version of Qwen 1.5 72B, is selected to add an intermediary model.

All these models are used with versions that are quantized¹, a technique used to reduce inference time and memory footprint, the quantization chosen for each model is given in bits per weight (bpw). With their quantization, they require around 8GB, 12GB, and 44GB of memory respectively.

Every model is evaluated on local instances, served locally with an engine-based API backend, using TabbyAPI² based on ExLLamaV2¹ multiple GPUs and without automatic splitting, using a workstation equipped with a Ryzen 9 7950x, 96GB of DDR5 memory running at 5600mhz and 2 Nvidia RTX 4090, each with 24 GB dedicated memory, running Ubuntu 23.10.

Experiments are carried out beforehand on various uncontrolled scenarios to define LLM parameters. With the sole aim of reducing non-determinism from one cycle to the next, the final parameters modified from the default engine parameters are as follows:

- **max_tokens**, maximum number of tokens in output: 300
- **min_p**, minimum percentage value that a token must reach to be considered (Value is scaled based maximum token probability): 0.05
- **temperature**, parameter that regulates the randomness: 0.2

The RAG is implemented using Langchain[Topsakal and Akinci, 2023] with an inference engine from HuggingFace³,

¹ <https://github.com/turboderp/exllamav2>

² <https://github.com/theroyallab/tabbyAPI>

³ <https://github.com/huggingface/text-embeddings-inference>

to locally execute an embedding model: BAAI/bge-large-en-v1.5[Xiao *et al.*, 2023], and an Elasticsearch⁴ local instance is used as vector database, both instances running on CPU and associated memory.

To evaluate system performance, each scenario is executed 10 times with each prompting style, and a grade is given to each response. Results then count the total number of points for each prompting category in general, and also for each defined metric associated with each question. The system’s complete processing time is also measured, in order to estimate the average latency of the different prompting styles.

The theoretical random action grade for each scenario is calculated as a baseline, using the following formula 1.

$$\text{grade}_s = \frac{\text{number_of_actions_rated_1}_s + 2 \times \text{number_of_actions_rated_2}_s}{\text{number_of_actions}_s} \quad (1)$$

Figure 2 shows an example of a scenario data representation using proposed natural language textual representation.

```
Current State of the House:
User 1 is in the Livingroom.
User is watching TV.
Previously: User was currently looking at TV

Livingroom: Curtains are Closed.
Lights: main, floor lamp are respectively Off, Off.
There is a TV in the room and its state is on.
CO2 level in room is 513ppm.
.....
Kitchen: Curtains are Closed.
Lights: main are Off.
CO2 level in room is 473ppm.

House was cleaned today.
Expected cleaning one time a week.
Centralized HVAC system is on with objective to 20°C.
Entrance smart Door is locked.
Time: 10:21 PM
Global house temperature is 20°C,
outside temperature is 5°C.
```

Figure 2: Extract of natural data representation on the scenario "Out of bed at night"

As previously mentioned, several data types have been included in the representations: numerical values, boolean values, and character strings. Using an LLM to process context allows the system to support all these data types.

Regarding the action list building algorithm, in one of the scenarios for example, it reduces the set of actions from more than 18 to 6 actions, which represents a 3-fold reduced set in this case. This reduced set takes into account the fact that lighting appliances are basic switches and cannot be set with different luminosity levels.

This limited set of actions makes it interesting to use a model randomly selecting one action from the set as a baseline. With a reduced action space, a random answer may not be so bad.

⁴<https://github.com/elastic/elasticsearch>

5 Results

5.1 Data representation

The first analysis of the results focuses on the advantage of using a natural representation versus a JSON representation.

Model	JSON	Natural
Qwen 1.5 72B	1.25	1.18
Qwen 1.5 14B	0.63	0.99
Starling Alpha 7B	1.03	1.18

Table 2: Comparing the two contextual representations: average grade of models per execution per scenario

Table 2 shows an average difference between results using JSON representation, or the more natural representation. It shows that the larger models are almost stable independently of the contextual representations, and on average even slightly better using JSON representation, by around 5.9%. However, on smaller models, the natural language representation greatly improves performances, with a 14.6% increase in average performance using Starling 7B, and a 57.1% increase using Qwen 1.5 14B. In terms of processing time, results are quite similar on average for models with both representations. On average, across all models, it leads to an increase in performance of 21.9%, despite the results on the larger Qwen model, making the natural representation more efficient.

5.2 User preferences

Figure 3 depicts the average grades obtained by the 3 chosen models on their responses to scenarios with the 4 different prompting styles and the 2 distinct representation types.

It first shows that Qwen 72B model is much more stable than the other two and that the prompting style does not have as much impact on response quality. Regarding the two smaller models, larger inconsistency in results can be noted as varying with the chosen prompting styles, particularly with Qwen 14B model.

The results of LLMs are compared with a baseline corresponding to the random choice of an action. The results remain on average behind any model without preferences (37.1% below the worst result with the "direct" prompting style). However, thanks to the algorithm reducing the set of actions and the fact that multiple responses are acceptable on each scenario, the random baseline obtains grades that are sometimes better than the ones of some experiments with LLMs.

The sequence of multiple questions requires the model to be consistent and to respect the expected instruction format. Furthermore, given that only prompt engineering is used to ensure the format, some prompting styles with multiple questions may lead to invalid responses, forcing the model to take a default action in the proposed setting. This default action is set to do nothing and to inform the user that it has failed to act. This reduces the performance of some models with some prompting styles. In figure 3 failure ratio measures the ratio of invalid responses. Qwen 14B results are especially below this baseline because they failed to answer in a large number of scenarios




Figure 3: Average grades by model, data representation and prompting styles

On average, "directPref" prompt models achieved a gain of 11.3% over basic prompt with natural representation, and even 20.0% with JSON representation.

The best results are achieved with JSON representation for "OpenQuestion" and "ThreeQuestion", leading to a 28.6% improvement over the "direct" prompting style.

With the natural representation, Qwen 14B model gives almost stable results with all prompting styles except "OpenQuestion", and the main difference with "OpenQuestion" seems to be linked to a high failure rate, as shown in figure 3. With JSON representation, failures are so high with all advanced prompting styles that the best results are obtained with the most basic prompt. In this case, "direct" prompting style is 30.7% more efficient than the "directPref" prompting style.

On Starling 7B, with both representations, the best results are obtained with "directPref", on the Starling representation, its results are 52.3% better than the direct representation and respectively 8.1% and 6.3% better than "OpenQuestion" and "ThreeQuestion" prompting styles.

In addition, the average processing time is reduced by 35.6% on Starling 7B Alpha using the "directPref" prompting style instead of "direct".

With JSON representation, the averaged grade results follow the same trend, but with lower overall values.

Table 3 shows that RAG-based prompting models can lead to better results than a single prompt containing all the data, as it is visible Qwen 72B. However, to date, LLM inference is still slow, and the use of complex prompts leads to a loss of accuracy in the test scenarios with smaller models.

Adding inference time to the balance with table 3 highlights Starling 7B Alpha results with "directPref". It outperforms almost all others except Qwen 72B with JSON repre-

Model	Prompting Style	Average grade	Proces. time (s)
Qwen 72	direct	1,06	9,04
	directPref	1,22	7,01
	OpenQuestion	1,18	24,43
	ThreeQuestion	1,26	42,13
Qwen 72 JSON	direct	1,05	9,02
	directPref	1,26	6,71
	OpenQuestion	1,35	25,20
	ThreeQuestion	1,35	41,47
Qwen 14	direct	1,09	1,22
	directPref	1,09	1,16
	OpenQuestion	0,65	8,83
	ThreeQuestion	1,11	15,54
Qwen 14 JSON	direct	0,98	1,28
	directPref	0,75	1,21
	OpenQuestion	0,47	9,39
	ThreeQuestion	0,32	16,62
Starling	direct	0,88	0,73
	directPref	1,34	0,47
	OpenQuestion	1,24	3,98
	ThreeQuestion	1,26	6,23
Starling JSON	direct	0,85	0,48
	directPref	1,15	0,48
	OpenQuestion	1,08	3,77
	ThreeQuestion	1,02	6,20

Table 3: Comparison of average grades for each prompting style and models with inference time

sentation and prompt chaining, but is 53.6 times faster, with a lower grade of just 0.7%.

With Qwen 72B, the use of more complex models (OpenQuestion, ThreeQuestion) can lead to better results, as seen previously in particular with JSON representation, but this comes with a trade-off: inference time. Qwen 72B model is already much slower due to its number of parameters, and due to the number of operations required to produce a single token. For instance, an inference with "directPref" takes 6.86 seconds on average (JSON and textual representation), whereas using "OpenQuestion" (which is around 60% faster than "ThreeQuestion") is 3.6 times slower.

In the current state of this type of hardware, it is impossible

to consider them as a viable alternative for managing a smart home automation system, with such reaction times.

6 Discussion

This section mainly discusses the results of the LLMs and prompting techniques chosen in order to make choices for real experiments. The advantages of this study are highlighted as well as the new challenges that are raised.

As seen previously, using larger LLM such as Qwen 72B allows greater stability in preference-free scenarios. Indeed, Qwen 72 is 19.9% better than Starling 7B in this case. However there are drawbacks, the first being the inference time as mentioned, and the second being the hardware infrastructure required. The quantized version of Qwen 1.5 72B requires 44GB of memory compared with around 8GB for Starling 7B Alpha with lower quantization.

Compared with Qwen 72B model using JSON representation, Starling 7B Alpha with "directPref" takes advantage of natural representation and achieves almost similar performances with much more complex prompting techniques. This makes the approach of using Starling 7B Alpha with this prompting style a good choice for future work. It gives similar performances concerning grades, and has a relatively low inference time (Average: 0.47s).

In addition, compared to Qwen 1.5 72B with natural representation and no preferences, Starling 7B Alpha's performance is 26.4% better using "directPref", with a processing time almost 20 times faster.

The results show the advantage of adhering to preferences. Drawbacks appear, however, with the additional average computation time for "OpenQuestion" and "ThreeQuestion" prompting styles, which use RAG. Using RAG brings no advantage in most cases. This is certainly due to the relatively small database. If the system required a larger database of preferences and rules, the results might have been different as it would not have been possible to give them directly through "directPref" prompting style. Based on current results, the best choice for a use case with a larger database would remain Starling 7B Alpha, with "OpenQuestion" prompting style and natural representation, as it provides results that are aligned with users preferences, with only 8.1% less average grade than Qwen 1.5 72B, while keeping an acceptable average inference time (3.98 seconds vs 25.20 seconds).

The use of smaller models poses the challenge of enforcing the output format. New methods have recently been proposed such as Outlines⁵, to strengthen the output grammar. They should be tested soon in order to analyze the cost in inference time and the improvement of the results of this paper.

Another approach is fine-tuning. The limits of using fine-tuning for user preferences were mentioned previously. However, fine-tuning on smart home domain data could allow the system to ensure the output format while improving the analysis of contextual data.

Although the proposed system is not comparable to other works due to its completely new approach, it has the advantages of its nature, as well as drawbacks. Firstly, it is not deterministic. Each configuration has been run 10 times, but some results are different each time, which means that the results of the system cannot be certified and a safety layer must be developed to secure some actions. Secondly, LLM queries are slow and require significant computing and memory resources: adaptability without training comes at a cost.

7 Conclusion

This paper presents a new architecture for a smart home automation system, using LLMs with user preferences to enhance personalised user experiences. This approach leverages the general knowledge provided by LLMs and combines it with naturally written rules and preferences to make contextually relevant decisions in line with user preferences. This architecture is proactive, able to adapt to any change in the environment thanks to the robustness provided by LLMs.

The user-centred action list builder takes advantage of this ability to reduce the set of actions at each step, as the representation of the environment also takes advantage of the name given to devices by the user to better support them.

The experimental results demonstrate the potential of this architecture to improve alignment with user preferences compared with an implementation without user preferences, showing up to 52.2% performance increase.

The study showed that, particularly with small models, using a natural representation instead of a JSON representation leads to an increase in performance, with an average 21.9% increase.

Although the system shows promising results on a set of defined scenarios, it also presents challenges due to stochastic behaviour and a slower inference time compared to traditional machine learning methods. These drawbacks are offset by the system's ability to adapt dynamically - without retraining - to changes in preferences, appliances and home configuration.

Future work will focus on implementing the system in a real-world smart home middleware system such as OpenHAB[Portalés *et al.*, 2019] to evaluate its performance with real users. Mechanisms will be proposed to allow users to naturally add and remove preferences, as well as to explore the automatic evolution of these rules and preferences.

⁵<https://github.com/outlines-dev/outlines>

Ethical Statement

There are no ethical issues.

References

- [Bai *et al.*, 2023] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- [Dey, 2001] Anind K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing*, 5(1):4–7, February 2001.
- [King *et al.*, 2023a] Evan King, Haoxiang Yu, Sangsu Lee, and Christine Julien. "Get ready for a party": Exploring smarter smart spaces with help from large language models. *Preprint*, 2023. Publisher: arXiv Version Number: 1.
- [King *et al.*, 2023b] Evan King, Haoxiang Yu, Sangsu Lee, and Christine Julien. Sasha: creative goal-oriented reasoning in smart homes with large language models. *Preprint*, 2023. Publisher: arXiv Version Number: 1.
- [Lewis *et al.*, 2020] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [Oliveira *et al.*, 2022] Pedro Filipe Oliveira, Paulo Novais, and Paulo Matos. Adaptive environment system to manage comfort preferences and conflicts. In Ana I. Pereira, Andrej Košir, Florbela P. Fernandes, Maria F. Pacheco, João P. Teixeira, and Rui P. Lopes, editors, *Optimization, Learning Algorithms and Applications*, pages 685–700, Cham, 2022. Springer International Publishing.
- [Peng *et al.*, 2020] Z. Peng, X. Li, and F. Yan. An adaptive deep learning model for smart home autonomous system. In *2020 International Conference on Intelligent Transportation, Big Data and Smart City (ICITBS)*, pages 707–710, Los Alamitos, CA, USA, jan 2020. IEEE Computer Society.
- [Portalés *et al.*, 2019] Cristina Portalés, Sergio Casas, and Kai Kreuzer. Challenges and trends in home automation: addressing the interoperability problem with the open-source platform openhab. *Harnessing the Internet of Everything (IoE) for Accelerated Innovation Opportunities*, pages 148–174, 2019.
- [Radford *et al.*, 2019] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- [Rashidi and Cook, 2009] P. Rashidi and D.J. Cook. Keeping the Resident in the Loop: Adapting the Smart Home to the User. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 39(5):949–959, September 2009.
- [Ruiz *et al.*, 2021] Ana Rubio Ruiz, David Villa Alises, Rubén Cantarero Navarro, María J. Santofimia Romero, Javier Dorado Chaparro, and Juan C. López López. Leveraging commonsense reasoning towards a smarter Smart Home. *Procedia Computer Science*, 192:666–675, 2021.
- [Shuvo, 2022] Salman Sadiq Shuvo. Home Energy Recommendation System (HERS): A Deep Reinforcement Learning Method Based on Residents' Feedback and Activity. *IEEE TRANSACTIONS ON SMART GRID*, 13(4), July 2022.
- [Topsakal and Akinci, 2023] Oguzhan Topsakal and Tahir Cetin Akinci. Creating large language model applications utilizing langchain: A primer on developing llm apps fast. In *International Conference on Applied Engineering and Natural Sciences*, volume 1, pages 1050–1056, 2023.
- [Weiser, 1991] Mark Weiser. The computer for the 21st century. *Scientific American*, 265(3):66–75, January 1991.
- [Xiao *et al.*, 2023] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to advance general chinese embedding, 2023.
- [Zhang *et al.*, 2023] Zhiyao Zhang, Yongxin Su, Mao Tan, and Rui Cao. Fusing domain knowledge and reinforcement learning for home integrated demand response online optimization. *Engineering Applications of Artificial Intelligence*, 121:105995, 2023.
- [Zhu *et al.*, 2023] Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. Starling-7b: Improving llm helpfulness and harmlessness with rlaif, November 2023.