PrivacyOracle: Configuring Sensor Privacy Firewalls with Large Language Models in Smart Built Environments

Brian Wang
wangbri1@g.ucla.edu
University of California, Los Angeles

Luis Antonio Garcia
la.garcia@utah.edu
Kahlert School of Computing
University of Utah

Mani Srivastava
mbs@ucla.edu
University of California, Los Angeles
and, Amazon*

2024 IEEE Security and Privacy Workshops (SPW) | 979-8-3503-5487-4/24/$31.00 ©2024 IEEE | DOI: 10.1109/SPW63631.2024.00028

***Abstract*—Modern smart buildings and environments rely on sensory infrastructure to capture and process information about their inhabitants. However, it remains challenging to ensure that this infrastructure complies with privacy norms, preferences, and regulations; individuals occupying smart environments are often occupied with their tasks, lack awareness of the surrounding sensing mechanisms, and are non-technical experts. This problem is only exacerbated by the increasing number of sensors being deployed in these environments, as well as services seeking to use their sensory data. As a result, individuals face an unmanageable number of privacy decisions, preventing them from effectively behaving as their own "privacy firewall" for filtering and managing the multitude of personal information flows. These decisions often require qualitative reasoning over privacy regulations, understanding privacy-sensitive contexts, and applying various privacy transformations when necessary. We propose the use of Large Language Models (LLMs), which have demonstrated qualitative reasoning over social/legal norms, sensory data, and program synthesis, all of which are necessary for privacy firewalls. We present *PrivacyOracle*, a prototype system for configuring privacy firewalls on behalf of users using LLMs, enabling automated privacy decisions in smart built environments. Our evaluation shows that *PrivacyOracle* achieves up to 98% accuracy in identifying privacy-sensitive states from sensor data, and demonstrates 75% accuracy in measuring social acceptability of information flows.**

***Index Terms*—Large Language Models, Privacy, Contextual Integrity, Smart Environments**

I. Introduction

Our physical spaces are progressively outfitted with a growing number of devices, with modern smart-built environments collecting and performing inferences over an increasingly rich and diverse set of sensory data. However, the sensing infrastructure faces novel privacy risks as interactions between devices and owners evolve. The amount of personal data captured in smart environments only continues to grow as the number and types of deployed devices increase over time. This is further exacerbated by advances in machine learning algorithms which have introduced a rapidly growing list of invasive inferences from a variety of sensor modalities. At the same time, individuals continue to visit new smart environments where control and awareness of sensory infrastructure range from complete ownership to non-existent. The result is a largely unregulated sensing environment where privacy decisions become unmanageable due to the scale of devices and novel yet invasive inferences over personal sensory data.

While there have been several types of systems proposed to alleviate the burden of privacy decisions and regulate the processing of sensory data, they often fall short of the requirements for modern smart environments. Approaches that learn and apply user privacy preferences require bespoke machine learning models based on a curated dataset of privacy perceptions which become outdated over time, limiting their effectiveness as users encounter new types of smart environments with unique applications and devices. Mechanisms for simplifying consent mechanisms may ameliorate the complexities of understanding privacy risks, but they do not mitigate user effort, thus making them ineffective in environments where awareness of sensory infrastructure is limited, and non-scalable when the set of applications and devices increase. Lastly, approaches to building trusted sources of access control and processing of sensory data still depend on user specification of privacy rules, falling short of the automated and informed privacy decision-making required of smart environments. Thus, it is necessary to have a mechanism acting as a firewall for the automatic filtering and management of personal sensory data, which we describe as *privacy firewalls*. These firewalls mediate the flow of information between sensing infrastructure and service providers, who may use personal data captured from data subjects to provide various services. This is shown in Figure 1.

Effective privacy firewalls in modern smart environments require several properties. First, they must be capable of qualitative reasoning about data-sharing decisions that involve utility, societal values, social pressures, as well as cultural and legal rules. This is especially important in the case of regulations, as they are costly to enforce. For example, the California Consumer Privacy Act (CCPA) is estimated to cost roughly 78 billion USD annually total imposed cost on businesses seeking to comply [1]. Second, they must be able to infer different

*Mani Srivastava holds concurrent appointments as a Professor of ECE and CS (joint) at UCLA and as an Amazon Scholar. This paper describes work performed at UCLA and is not associated with Amazon.

DOI 10.1109/SPW63631.2024.00028

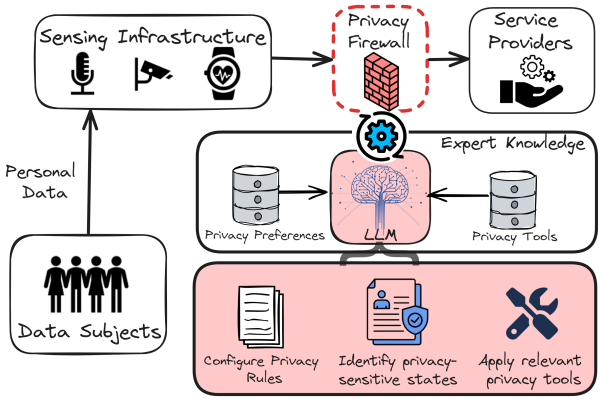

Fig. 1. High-level overview of privacy firewalls, where expert knowledge in the form of LLMs can be used to configure sharing of sensory data.

states of users in the environment from low-level sensor data, and whether those states present privacy risks. For example, a user who is lying down at a particular place and time may refer to a private event of sleeping. With an understanding of the context surrounding sensory information, privacy firewalls can effectively operate independently of users. Third, a privacy firewall should be able to quickly adapt to different tools for preserving privacy. While many ML algorithms are uncovering privacy risks in data, a variety of techniques are emerging that can be utilized to preserve privacy, particularly under the umbrella of generative AI [2], [3], enabling various tools to be employed for privacy-sensitive data under different conditions.

Recently, Large Language Models such as ChatGPT [4] have demonstrated tremendous capability in breaking down high-level tasks into several technical steps [5], as well as being able to summarize and provide question answering, such as for privacy policies [6]. In addition, LLMs have passed rigorous exams in law [7] and medicine [8], suggesting they possess deep knowledge of legal, social, and ethical norms. However, it has yet to be studied how much latent knowledge about privacy norms is captured within LLMs, and how they can be applied in configuring privacy firewalls.

We hypothesize that LLMs can be harnessed to accomplish the tasks necessary for effective privacy firewalls in smart environments. In particular, they are capable of automated privacy decision-making and processing, a task that would normally involve significant user involvement or ML models trained on privacy preferences. Neither are appropriate for modern smart environments given the scale of devices present and the necessary adaptation to new regulations and norms. Importantly, LLMs are capable of qualitative reasoning over data processing decisions due to possessing the significant degree of world knowledge required in making these decisions.

The main contribution of this work, *PrivacyOracle* [1], is an LLM-based system for managing and configuring privacy firewalls in smart environments. We investigate the performance of the privacy firewall on several tasks, as shown in Figure 1: validation of informational flows under particular legal rules and social norms, inference of privacy-sensitive states from low-level sensor data from the environment, and lastly selection of relevant data transformations under different privacy and utility requirements.

[1]https://github.com/nesl/PrivacyOracle

II. Related Work

In this section, we cover various existing ideas for automating privacy decisions and managing flows of data.

A. Automating Privacy Decisions

The idea of privacy firewalls and controlling flows of sensory data has also been suggested in the past. Early work on privacy-aware data streaming involved personal data stores, providing a trusted source of storage and access control operating over user-specified policies [9], [10], with similar ideas proposed today in the form of hub-based data flow control [11]–[13]. However, these approaches still require a significant level of user effort (such as for specifying policies), cannot qualitatively reason about regulations and norms, and are unable to selectively choose different privacy tools based on current privacy and utility requirements.

Other works aim to help manage the burden of privacy decisions in smart environments. Although the concept of privacy assistants is over a decade old [14], systems have been developed specifically for both mobile and IoT settings [15]–[17] which involve studying user preferences and training ML models to predict privacy preferences of new users. As a result, it allows automatic configuration of these settings given new application environments without user involvement. However, these systems typically lack knowledge of social norms and legal regulations for privacy, and can not easily integrate new knowledge of sensor risks.

There has been some prior work that seeks to encode legal rules and smart home privacy preferences into the format of contextual integrity [18]–[20], which allows automated reasoning over informational flows. However, these works cannot validate the appropriateness of new informational flows that may arise with adding new sensors and services in an environment. More specifically, these systems cannot generalize to new scenarios and reason about their privacy behaviors.

B. Reducing Effort in Privacy Decisions

Several approaches aim to create more intuitive representations of sensing information in an environment, effectively reducing the user effort in making privacy decisions in how their data is used and shared [6], [21]. At the same time, several works employ language models to assist in parsing and understanding privacy policy documents, such as [22], [23]. Both of these approaches aim to create more usable privacy systems. However, these works do not address the more fundamental challenge of automatically identifying and handling privacy violations using various tools without involving the user.

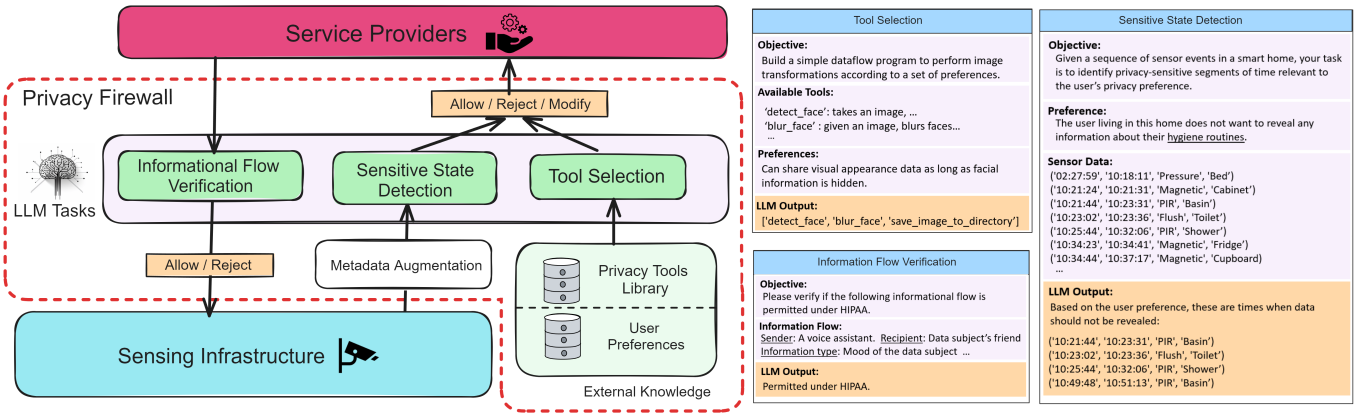

Fig. 2. Architecture of *PrivacyOracle*, our privacy firewall for regulating the flow of sensory data in smart-built environments via several LLM tasks.

C. Configuring Systems using LLMs

Although LLMs haven't been deeply investigated in applying privacy controls over sensory data, the general idea of applying LLMs to configure systems and manage data has been explored before. Several works aim to apply LLMs in configuring routers [24], [25], configuration tuning for distributed systems [26], and detecting vulnerabilities in software [27], [28]. Other LLM-based systems have also been created that are capable of operating over sensor data and environmental events [29]–[31].

III. System Design

The main objective of this work is to configure a *privacy firewall*, which regulates informational flows between sensing infrastructure and smart environments. However, to limit the scope of this work, we consider smart environments which consist of a sensing infrastructure that is willing to cooperate with data subjects' privacy preferences. In addition, we assume that data subjects share a similar set of privacy requirements, and do not consider conflicting preferences. In this work, we consider examples like smart homes, where the primary data subject is a homeowner who wishes to control information flowing from a personally managed set of devices. However, in section V we discuss other settings as well, including conflicts in privacy requirements and adversarial scenarios.

A privacy firewall may be co-located on the same premises as the sensing infrastructure (such as a trusted hub [12], [13]) or may exist as a separate storage/compute platform on an external network (such as personal data vaults [9] or databoxes [10]). As shown in Figure 1, the sensing infrastructure captures the personal information of data subjects who enter a smart-built environment (possibly with or without their consent). The privacy firewall then acts as a filter between information sent between the sensing infrastructure and some 3rd party service providers to ensure the privacy of data subjects.

In this work, we investigate the necessary services for privacy firewalls in protecting sensory data. These services are shown in Figure 2 as informational flow verification, sensitive state detection, and tool selection. Each service is powered by LLMs, which perform reasoning over unstructured privacy requirements, as well as qualitative reasoning over latent knowledge of privacy norms and regulations. We will now describe each service and how it interacts with the flow of information, shown in Figure 2.

A. Informational flow verification

The first service is to verify information flow requests from a service provider to a sensing infrastructure. This service must identify the necessary knowledge to configure the privacy rules for sharing these informational flows.

We use the same definition of informational flows as described in Contextual Integrity (CI) [32]. Contextual Integrity is a theory of privacy based on recognizing and evaluating appropriate flows of personal information in the absence of effective consent mechanisms. Privacy is achieved when a flow of personal information is deemed appropriate by entrenched social and legal norms - in our experiments, we evaluate appropriateness as compliance with privacy regulations (i.e. HIPAA) and social norms. Under CI, a personal informational flow is described by five parameters: data subject (who the data is about), the sender (the entity that transmits the data), the recipient (the entity that receives the data), information type (describes the category of the information), and lastly, the transmission principle (applies constraints under which the information flows). Thus, a personal information flow can be described with a tuple of 5 elements: (data subject, sender, recipient, attributes, and transmission principles).

As shown in Figure 2, an informational flow is proposed by a service provider, and evaluated by the LLM. The evaluation should consider both privacy regulations (such as HIPAA) and norms of society. An example of this reasoning is shown on the right side of Figure 2, under "Informational Flow Verification". This task involves specifying the informational flow parameters under a particular legal or social context and asking the LLM to decide if the flow is acceptable. *PrivacyOracle* may then use the response generated by the LLM to decide if an information flow is acceptable or not, and share information accordingly.

B. Sensitive State Detection

The second service in a privacy firewall is to identify sensitive segments of sensor data that should be hidden based on privacy preferences. This service requires sensor data to be augmented with metadata information (such as sensor type, location, and names), that grants additional context for identifying the sensitivity of data. This task filters out sensor data based on a privacy preference expressed in natural language.

An example of this filtering process is shown on the right side of Figure 2, shown under "Sensitive State Detection". We provide a preference (such as hiding sensor data relating to hygiene activities) and give the LLM a sequence of sensor data augmented with certain metadata to assist in identifying privacy-sensitive segments. The response from the LLM may then be used to filter out sensor data flowing to the service provider based on the time intervals obtained from the LLM.

C. Tool Selection

The last service of a privacy firewall is to identify the appropriate tools for transforming sensor data into a format that is acceptable from a privacy perspective. This allows privacy firewalls to not only accept or reject informational flows but modify them such that they are acceptable. Tool selection involves reasoning over both the privacy preferences of users and a library of tools, which allows automated selection of a particular tool given different scenarios. Furthermore, this service not only selects tools but can generate dataflow pipelines using other preprocessing tools as well, which improves the interoperability of different tools without having to worry about modifying the tool interfaces.

Figure 2 describes an example of this pipeline generation process, where a set of available tools and their high-level descriptions are given to an LLM, in addition to a privacy preference. The LLM must then perform qualitative reasoning over each different tool (which may offer various privacy-utility tradeoffs), and identify the correct tool and pipeline in which to execute that tool. *PrivacyOracle* then parses the LLM output, generates the pipeline, and executes it on sensor data to obtain different privacy and utility metrics.

IV. Evaluation

This section describes experiments using LLMs in accomplishing various tasks for privacy firewalls. In our experiments, we use GPT-3.5 and GPT-4.0 as our LLM [2].

Validating information flows with privacy regulations

Setup. We choose HIPAA as a case study for evaluating how well LLMs can validate a hypothesized informational flow against its latent knowledge of privacy regulations. We manually create 16 different flows of information using the 5 parameters of CI. 8 of the 16 flows are legally acceptable under the HIPAA Privacy Rule (manually verified by the authors), while the other 8 are not. Similarly, 8 of the 16 flows involve a data subject which is a 1st party, while the other 8 involve a data subject which is a 3rd party. The goal of this experiment is to identify the LLM's capability in distinguishing the necessary characteristics for a flow to be acceptable under a 1st party vs. a 3rd party data subject. More specifically, a flow is valid depending on the entity providing consent, and we seek to establish if an informational flow is valid depending on the data subject (1st vs. 3rd party) and which entity has provided consent to the informational flow. This requires an LLM to understand which party consent must be obtained from under HIPAA. We use GPT-3.5 as our LLM.

Results. LLMs appear to have decent success at identifying which entities are critical for providing consent in 1st and 3rd party data subjects, with a low false positive rate of 6.25% and a false negative rate of 6.25%.

Measuring Social Acceptability of Informational Flows

In addition to measuring how LLMs can validate informational flows against privacy regulations, we also evaluate how well they can reason over privacy norms.

Setup. We obtain 144 different information flows generated from a subset of CI parameters described in [19]. In addition to the 4 options for data recipients and transmission principles shown in Figure 3, we provide 3 options for the sender parameter (sleep monitor, fitness tracker, and door lock), and 3 options for the data type (location, audio, and exercise routine). The data subject is, by default, the owner of the device. We then query the LLM to obtain acceptability scores between -1.5 and 1.5, which aims to measure how acceptable an informational flow is given the varying parameters, with -1.5 being the least acceptable and 1.5 being the most acceptable. The goal of this experiment is to study how much the acceptability scores obtained from an LLM reflect the acceptability scores from previous user studies on the same informational flows. We use GPT-3.5 in this experiment.

Results. Our results are shown in Figure 3, which aims to measure the difference in acceptability scores obtained from the LLM and a previous user study on the same informational flows [19]. We average the acceptability scores obtained from the LLM for each flow which shares the same data recipient and transmission principle, resulting in a matrix of 16 different flows. We then measure the absolute difference between the acceptability score from the user study and the average acceptability obtained from the LLM. We assign a sign to the resulting value based on whether the scores match in terms of their sign. For example, if the score of the user study is acceptable (ranges from 0 to 1.5) and the score obtained from the LLM is unacceptable (ranges from -1.5 to 0), then the assigned sign is negative. While the differences between the LLM and user acceptability scores are quite significant, the majority of informational flows match in terms of binary acceptability (roughly 75%), demonstrating some agreement between the social norms of privacy expressed by users and knowledge of social norms by the LLM.

[2]Some tasks can be performed adequately using GPT-3.5, while others required the more capable GPT-4.0. However, the cost of each differs in an order of magnitude, requiring judicious use of both.

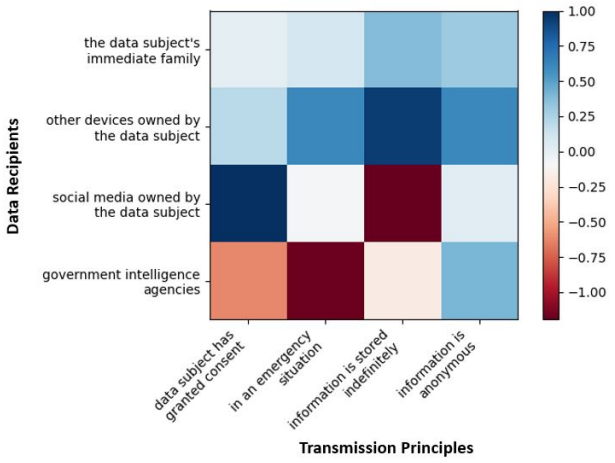

Fig. 3. Difference in acceptability scores of information flows with given recipient/transmission principle pairs, measured between LLM and a previous user study [19].

Identifying Privacy-Sensitive States from Sensor Data

To help privacy firewalls reason about the sensitivity of sensor data and whether they leak private information, we believe LLMs can provide sufficient background knowledge about sensors and their context to identify potentially sensitive segments of sensor data.

Setup. We use a dataset capturing 35 days of activities of daily living (ADLs) [33]. Each day contains data from IoT sensors and their corresponding ADL label. We augment each sensor data sample with information about what sensors are being activated (e.g. microwave activated), the time of activation, as well as the type of sensor (e.g. magnetic). We split this augmented sensor data into segments of 24 hours. We evaluate the performance of LLMs for identifying specified privacy-sensitive states from these segments of data. Table I describes 3 privacy sensitive states, which we ask the LLM to identify from the given segments of sensor data. For example, we asked the LLM to identify time intervals where sensor data revealed information about hygienic activities. The evaluation was performed using GPT-4 with a fixed seed and temperature of 0. The ground truth was obtained from the labeled activities of daily living, with each privacy-sensitive state associated with several categories of activities. For hygienic activities, we associated them with labels of "Toileting", "Grooming", and "Showering". For a sedentary lifestyle, we associated it with the labels "Spare_Time/TV" and "Sleeping". Lastly, for house occupancy, we associated it with the label "Leaving".

Results. We measure the mean average error (MAE), intersection-over-union (IoU), and F1 score in the time intervals predicted by the LLM as privacy-sensitive and the ground truth associated labels. We find that the poor MAE performance is mainly due to outliers where large time gaps may occur between one privacy-sensitive state and another. A missed state would result in a significant time difference until the next associated state. However, the IoU and F1 scores demonstrate that there is significant overlap and agreement on what the LLM identifies as relevant privacy-sensitive states.

Privacy-Sensitive State	MAE (seconds)	IoU	F1
Hygenic Activities	307.04	0.684	0.831
Sedentary Lifestyle	865.31	0.844	0.983
House Occupancy	294.84	0.961	0.701

TABLE I
DETECTION ACCURACY OF GPT-4 FOR PRIVACY-SENSITIVE STATES FROM SENSOR DATA.

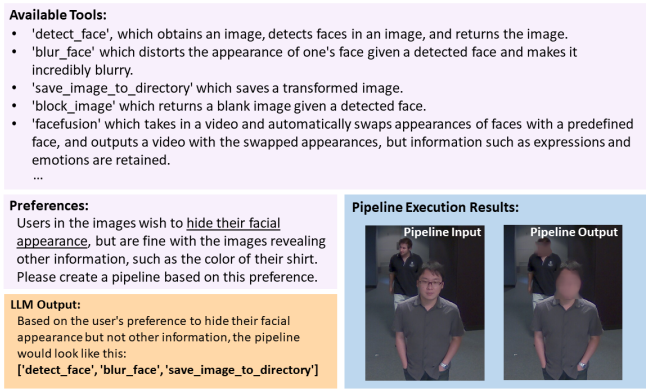

Fig. 4. Response examples of GPT-4 for adapting privacy tools to natural language requirements

Automatic Adaptation of Privacy Tools to Natural Language Requirements

While our previous experiments demonstrate that LLMs are capable of verifying whether an informational flow violates privacy. In this experiment, we take this idea a step further to identify what kind of transformations we may apply to a flow of sensory data to ensure privacy.

Setup. We perform a set of experiments on the Chokepoint person identification dataset [34]. This dataset consists of multiple video sequences recording subjects from 2 different doorway cameras. We also use a set of different privacy requirements over a stream of images, shown in Table II.

- Requirement A: "Users in the images wish to hide any trace of them being part of the video."
- Requirement B: "Users in the images wish to hide their facial appearance, but are fine with the images revealing other information, such as the color of their shirt."
- Requirement C: "Users in the images wish to hide their facial appearance, and they would like to enable emotion classification services on their facial data."

We use GPT-4 to generate a simple dataflow program that uses a library of different privacy tools while adhering to the set of natural language privacy requirements. In addition to some data loading and preprocessing functions, we use 3 privacy tools: A **Blur faces** tool applies a Gaussian blur over a detected face. A **Block images** masks the entire image when a face is detected. Lastly, a **FaceFusion** tool [2] performs high-fidelity face-swapping, where faces in a sequence of images can be replaced by a source face image while also retaining features such as environmental lighting and expressions. The resulting programs for each requirement is as follows:

- Requirement A: ['detect face', 'block image', 'save image to directory']

Requirement	Age F1	Gender F1	Race F1	Emotion F1
A	0.208	0.373	0.151	0.081
B	0.0	0.0	0.0	0.0
C	0.187	0.371	0.288	0.407

TABLE II
PRIVACY (AGE, GENDER, RACE) AND UTILITY SCORES (EMOTION) OF PIPELINES GENERATED BY DIFFERENT PRIVACY REQUIREMENTS.

- Requirement B: ['detect face', 'blur face', 'save image to directory']
- Requirement C: ['detect face', 'convert images to video', 'facefusion', 'convert video to images', 'save image to directory']

Note that the facefusion tool also operates over video. To test alternative pipelines we consider other processing tools (such as for converting images to video). We then execute these dataflow programs and evaluate the privacy and utility characteristics of the transformed data. We measure the age, gender, race (private) and emotion (utility) of different users via facial analysis provided by DeepFace [35].

Results. We executed the pipelines generated by the LLM, and measure the privacy properties of the resulting transformed images (age, gender, and race), as well as certain utility requirements that may be desirable from the user's perspective (emotion). Table II shows the F1 scores of various facial characteristics. Depending on the requirement and executed program from GPT-4, each tool creates different privacy-utility tradeoffs. The **Block images** tool provides a high degree of privacy (no information about age, gender, or race can be inferred), but incurs a cost when it comes to utility (emotion recognition). The **Blur faces** tool also yields decent privacy performance while possibly maintaining other types of utility information (e.g. occupancy detection). Lastly, the **FaceFusion** tool also yields good privacy performance but is also somewhat effective for the utility requirement (emotion detection).

V. DISCUSSION

A. Improving LLM outcomes

One of the main challenges in these experiments was obtaining good, yet reasonably consistent results. In the context of validating informational flows, we found that LLMs provided more consistent results when names were assigned to subjects and recipients, rather than defining them by their roles - in our experiments it appears that LLMs tend to conflate multiple roles with the same person unless explicitly specified. Another important discovery was requiring the LLM's response to be structured in a particular set of steps before providing an answer, such as identifying the data subject before generating the answer. This is similar to the idea of zero-shot chain of thought prompting, proposed in [36], which requests the LLM to 'think step by step'.

B. Generalization to other smart environments

Some of the assumptions made in this work may not be realistic or generalizable to all smart environments. We assume a cooperative sensing infrastructure operating a privacy firewall. However, not all smart-built environments will cooperate with their data subjects' privacy requirements. In addition, we have not considered cases of conflicting privacy requirements between data subjects. We leave these avenues of research for future work.

C. Relation to network security

The concept of a privacy firewall shares many characteristics with network firewalls and intrusion detection systems. Our task of validating incoming flows of information against privacy regulations and social norms is similar to standard packet filtering firewalls [37] for restricting the flow of network data based on certain policies. The task of identifying different privacy-sensitive patterns is akin to signature-based intrusion detection systems found in next-gen firewalls [38], [39]. Lastly, the task of adapting new privacy tools to process data is similar to the concept of 'active response' mechanisms [40] used as part of the deep packet inspection intrusion prevention systems [41]. Unlike network security mechanisms, a privacy firewall must be able to perform significant qualitative reasoning with effective use of world knowledge about privacy.

D. Ethical and social implications

Lastly, it is worth considering some of the ethical and societal implications of automated privacy regulation. Privacy, particularly in public spaces, has a tradeoff with security and safety, where efforts to manage crime and emergencies may be hindered when data is restricted. Managing this tradeoff is essential for future work. Automated regulation itself is a costly endeavor and may create biases due to socioeconomic disparities. Finally, much technological innovation comes from access to sensory data and it may be challenging to do research under automated privacy regulations.

VI. CONCLUSION

In this work we propose the idea of using LLMs to manage sensory information flows in smart built environments, thus acting as a privacy firewall. We study the performance of LLMs on three services necessary for privacy firewalls: validation of sensory informational flows against privacy regulations and norms, detection of privacy-sensitive states from low level sensory data, and selection of appropriate data transformation pipelines to preserve privacy. Our initial experiments demonstrate the potential of LLM-based privacy reasoning, and suggests that a world where informational flows are automatically configured for privacy regulations, norms and preferences may be closer than anticipated.

ACKNOWLEDGEMENTS

The research reported in this paper was sponsored in part by the the National Science Foundation (NSF) under awards 1705135 and 2124252, the NIH mDOT Center under Award #1P41EB028242, the DEVCOM ARL under Cooperative Agreement #W911NF-17-2-0196, and the DARPA ANSR Program under Contract #FA875023C0519. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies.

REFERENCES

[1] D. Castro, L. Dascoli, and G. Diebold, "The looming cost of a patchwork of state privacy laws," Information Technology and Innovation Foundation, Tech. Rep., 2022.

[2] "facefusion/facefusion," Feb. 2024, original-date: 2023-08-17T19:59:55Z. [Online]. Available: https://github.com/facefusion/facefusion

[3] Voice.ai, "RVC V2 Voice Models," Oct. 2023. [Online]. Available: https://voice.ai/hub/voices/rvc-v2-voice-models/

[4] "Introducing ChatGPT," 2022. [Online]. Available: https://openai.com/blog/chatgpt

[5] T. Gupta and A. Kembhavi, "Visual Programming: Compositional visual reasoning without training," Nov. 2022, arXiv:2211.11559 [cs]. [Online]. Available: http://arxiv.org/abs/2211.11559

[6] S. Pan, T. Hoang, D. Zhang, Z. Xing, X. Xu, Q. Lu, and M. Staples, "Toward the Cure of Privacy Policy Reading Phobia: Automated Generation of Privacy Nutrition Labels From Privacy Policies," Jun. 2023, arXiv:2306.10923 [cs]. [Online]. Available: http://arxiv.org/abs/2306.10923

[7] A. B. A. Journal, "Latest version of ChatGPT aces bar exam with score nearing 90th percentile," 2023. [Online]. Available: https://www.abajournal.com/web/article/latest-version-of-chatgpt-aces-the-bar-exam-with-score-in-90th-percentile

[8] "ChatGPT Passes US Medical Licensing Exam Without Clinician Input," 2023. [Online]. Available: https://healthitanalytics.com/news/chatgpt-passes-us-medical-licensing-exam-without-clinician-input

[9] M. Mun, S. Hao, N. Mishra, K. Shilton, J. Burke, D. Estrin, M. Hansen, and R. Govindan, "Personal data vaults: a locus of control for personal data streams," in *Proceedings of the 6th International COnference*, 2010, pp. 1–12.

[10] A. Chaudhry, J. Crowcroft, H. Howard, A. Madhavapeddy, R. Mortier, H. Haddadi, and D. McAuley, "Personal data: thinking inside the box," 2015.

[11] L. Wang, U. Khan, J. Near, Q. Pang, J. Subramanian, N. Somani, P. Gao, A. Low, and D. Song, "{PrivGuard}: Privacy regulation compliance made easier," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 3753–3770.

[12] H. Jin, G. Liu, D. Hwang, S. Kumar, Y. Agarwal, and J. I. Hong, "Peekaboo: A hub-based approach to enable transparency in data processing within smart homes," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 303–320.

[13] H. Chi, Q. Zeng, X. Du, and L. Luo, "Pfirewall: Semantics-aware customizable data flow control for home automation systems," *arXiv preprint arXiv:1910.07987*, 2019.

[14] M. Langheinrich, "A privacy awareness system for ubiquitous computing environments," in *UbiComp 2002: Ubiquitous Computing: 4th International Conference Göteborg, Sweden, September 29–October 1, 2002 Proceedings 4*. Springer, 2002, pp. 237–245.

[15] A. Das, M. Degeling, D. Smullen, and N. Sadeh, "Personalized privacy assistants for the internet of things: Providing users with notice and choice," *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 35–46, 2018.

[16] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti, "Follow my recommendations: A personalized privacy assistant for mobile app permissions," in *Twelfth symposium on usable privacy and security (SOUPS 2016)*, 2016, pp. 27–41.

[17] J. Colnago, Y. Feng, T. Palanivel, S. Pearman, M. Ung, A. Acquisti, L. F. Cranor, and N. Sadeh, "Informing the design of a personalized privacy assistant for the internet of things," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.

[18] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum, "Privacy and contextual integrity: Framework and applications," in *2006 IEEE symposium on security and privacy (S&P'06)*. IEEE, 2006, pp. 15–pp.

[19] N. Apthorpe, Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster, "Discovering smart home internet of things privacy norms using contextual integrity," *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, vol. 2, no. 2, pp. 1–23, 2018.

[20] N. Abdi, X. Zhan, K. M. Ramokapane, and J. Such, "Privacy Norms for Smart Home Personal Assistants," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI '21. New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–14. [Online]. Available: https://doi.org/10.1145/3411764.3445122

[21] P. Emami-Naeini, J. Dheenadhayalan, Y. Agarwal, and L. F. Cranor, "An informative security and privacy "nutrition" label for internet of things devices," *IEEE Security & Privacy*, vol. 20, no. 2, pp. 31–39, 2021.

[22] R. Khandelwal, T. Linden, H. Harkous, and K. Fawaz, "{PriSEC}: A privacy settings enforcement controller," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 465–482.

[23] T. A. Rahat, M. Long, and Y. Tian, "Is your policy compliant? a deep learning-based empirical study of privacy policies' compliance with gdpr," in *Proceedings of the 21st Workshop on Privacy in the Electronic Society*, 2022, pp. 89–102.

[24] R. Mondal, A. Tang, R. Beckett, T. Millstein, and G. Varghese, "What do llms need to synthesize correct router configurations?" in *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*, 2023, pp. 189–195.

[25] C. Wang, M. Scazzariello, A. Farshin, D. Kostic, and M. Chiesa, "Making network configuration human friendly," *arXiv preprint arXiv:2309.06342*, 2023.

[26] G. Somashekar and R. Kumar, "Enhancing the configuration tuning pipeline of large-scale distributed applications using large language models (idea paper)," in *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, 2023, pp. 39–44.

[27] D. Noever, "Can large language models find and fix vulnerable software?" *arXiv preprint arXiv:2308.10345*, 2023.

[28] K. Yang, J. Liu, J. Wu, C. Yang, Y. R. Fung, S. Li, Z. Huang, X. Cao, X. Wang, Y. Wang *et al.*, "If llm is the wizard, then code is the wand: A survey on how code empowers large language models to serve as intelligent agents," *arXiv preprint arXiv:2401.00812*, 2024.

[29] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg, "Progprompt: Generating situated robot task plans using large language models," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 523–11 530.

[30] H. Xu, L. Han, M. Li, and M. Srivastava, "Penetrative ai: Making llms comprehend the physical world," *arXiv preprint arXiv:2310.09605*, 2023.

[31] Z. Leng, A. Bhattacharjee, H. Rajasekhar, L. Zhang, E. Bruda, H. Kwon, and T. Plötz, "Imugpt 2.0: Language-based cross modality transfer for sensor-based human activity recognition," *arXiv preprint arXiv:2402.01049*, 2024.

[32] H. Nissenbaum, *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2020.

[33] F. Ordez, "Activities of Daily Living (ADLs) Recognition Using Binary Sensors," UCI Machine Learning Repository, 2013, DOI: https://doi.org/10.24432/C5J02M.

[34] Y. Wong, S. Chen, S. Mau, C. Sanderson, and B. C. Lovell, "Patch-based probabilistic image quality assessment for face selection and improved video-based face recognition," in *IEEE Biometrics Workshop, Computer Vision and Pattern Recognition (CVPR) Workshops*. IEEE, June 2011, pp. 81–88.

[35] S. I. Serengil and A. Ozpinar, "Hyperextended lightface: A facial attribute analysis framework," in *2021 International Conference on Engineering and Emerging Technologies (ICEET)*. IEEE, 2021, pp. 1–4. [Online]. Available: https://doi.org/10.1109/ICEET53442.2021.9659697

[36] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," *Advances in neural information processing systems*, vol. 35, pp. 22 199–22 213, 2022.

[37] S. M. Bellovin and W. R. Cheswick, "Network firewalls," *IEEE communications magazine*, vol. 32, no. 9, pp. 50–57, 1994.

[38] A. Khraisat, I. Gondal, P. Vamplew, and J. Kamruzzaman, "Survey of intrusion detection systems: techniques, datasets and challenges," *Cybersecurity*, vol. 2, no. 1, pp. 1–22, 2019.

[39] Z. Ahmad, A. Shahid Khan, C. Wai Shiang, J. Abdullah, and F. Ahmad, "Network intrusion detection system: A systematic study of machine learning and deep learning approaches," *Transactions on Emerging Telecommunications Technologies*, vol. 32, no. 1, p. e4150, 2021.

[40] B. Caswell, J. Beale, and A. Baker, *Snort intrusion detection and prevention toolkit*. Syngress, 2007.

[41] K. Scarfone, P. Mell *et al.*, "Guide to intrusion detection and prevention systems (idps)," *NIST special publication*, vol. 800, no. 2007, p. 94, 2007.