

VPVET: Vetting Privacy Policies of Virtual Reality Apps

Yuxia Zhan
dabeidouretriever@sjtu.edu.cn
Shanghai Jiao Tong University
Shanghai, China

Yan Meng*
yan_meng@sjtu.edu.cn
Shanghai Jiao Tong University
Shanghai, China

Lu Zhou
zhoulu@xidian.edu.cn
Xidian University
Xi'an, Shaanxi, China

Yichang Xiong
yxiong2@gmu.edu
George Mason University
Fairfax, Virginia, USA

Xiaokuan Zhang
xiaokuan@gmu.edu
George Mason University
Fairfax, Virginia, USA

Lichuan Ma
lcma@xidian.edu.cn
Xidian University
Xi'an, Shaanxi, China

Guoxing Chen
guoxingchen@sjtu.edu.cn
Shanghai Jiao Tong University
Shanghai, China

Qingqi Pei
qqpei@mail.xidian.edu.cn
Xidian University
Xi'an, Shaanxi, China

Haojin Zhu
zhu-hj@sjtu.edu.cn
Shanghai Jiao Tong University
Shanghai, China

ABSTRACT

Virtual reality (VR) apps can harvest a wider range of user data than web/mobile apps running on personal computers or smartphones. Existing law and privacy regulations emphasize that VR developers should inform users of what data are collected/used/shared (CUS) through privacy policies. However, privacy policies in the VR ecosystem are still in their early stages, and many developers fail to write appropriate privacy policies that comply with regulations and meet user expectations. In this paper, we propose VPVET to automatically vet privacy policy compliance issues for VR apps. VPVET first analyzes the availability and completeness of a VR privacy policy and then refines its analysis based on three key criteria: granularity, minimization, and consistency of CUS statements. Our study establishes the first and currently largest VR privacy policy dataset named VRPP, consisting of privacy policies of 11,923 different VR apps from 10 mainstream platforms. Our vetting results reveal severe privacy issues within the VR ecosystem, including the limited availability and poor quality of privacy policies, along with their coarse granularity, lack of adaptation to VR traits and the inconsistency between CUS statements in privacy policies and their actual behaviors. We open-source VPVET system along with our findings at repository <https://github.com/kalamoo/PPAudit>, aiming to raise awareness within the VR community and pave the way for further research in this field.

1 INTRODUCTION

Enhanced by multi-modal sensors and binocular visual rendering techniques, virtual reality (VR) offers an unparalleled immersive experience for users, establishing itself as the fundamental technology for the meta-verse. Major tech companies are heavily investing in VR headsets, with Meta releasing its most advanced Meta Quest Pro [33] and Apple developing their latest Apple Vision Pro [35]. As VR hardware continues to advance, there is a significant surge in the growth of VR applications (apps). Beyond its popularity in the gaming industry, VR apps are widely designed and deployed in various scenarios, including art [74], education [72], healthcare [50], tourism [76], military training [62], real estate [42],

retail [60], sports [46], and virtual meetings [56]. In August 2023, a VR company called Varjo closed a multi-million dollar deal to supply headsets for the US Army [62], showcasing the significance and promising future of VR technology.

Despite the widespread popularity of VR, researchers have found that the data collection/usage/sharing (CUS) process in VR apps faces an increasing risk of privacy leakage [10, 15, 18, 25, 31, 38, 54]. These privacy risks are more apparent in VR apps because, during the transition from deploying apps on desktop and traditional mobile devices (e.g., smartphones) to VR, the number and types of sensors and input/output (I/O) devices undergo significant increases, resulting in vast amounts of users' personal information being collected [2]. For instance, the measurements from cameras or infrared trackers of VR headsets and hand controllers can reveal users' biometric information (e.g., height and body shape), while data from eye-tracking sensors can expose users' opinions toward virtual content (e.g., user interest in advertising content [28]). Additionally, recent research indicates that users can be de-anonymized with over 90% accuracy among a pool of 50,000+ individuals based on just 100 seconds of motion data in VR games [55].

Privacy concerns related to VR apps are not only drawing attention from academia but are also emphasized with the promulgation of privacy laws and regulations such as Personal Information Protection Law (PIPL) in China [59], General Data Protection Regulation (GDPR) in European Union [27], and California Consumer Privacy Act (CCPA) in United States [20]. All of them stipulate that users have the *right to know* about the data usage process. Specifically, *privacy policies* serve as an important interface that allows users to understand how their personal data are collected, stored, and processed by any specific app in a *transparent* manner. Ideally, a well-written privacy policy can let users make informed decisions *before* using these apps based on whether their personal information is being handled appropriately.

Motivations. Unfortunately, prior studies indicate that 91% of users typically agree to privacy policies by simply clicking the checkbox, while skipping reading their contents [24, 43]. What's even worse, in the current VR ecosystem, our preliminary analysis (Section 3) reveals that for a significant proportion of VR apps, their

*Yan Meng is the corresponding author.

privacy policies neither meet the legal requirements nor satisfy user expectations, especially in the following three aspects. (1) *Poor accessibility*: Some VR apps' developers and the platforms where these app are published do not display the corresponding privacy policy. (2) *Lacking VR-specific content*: even if the privacy policy is accessible, it lacks details on how user data is handled in VR scenarios. (3) *Vagueness and misrepresentation*: the privacy policy either uses coarse descriptions or excessively claims to collect users' data, which deviates from apps' actual practices. The above observations motivate us to vet the compliance of privacy policies in the whole VR ecosystem.

Challenges. However, vetting privacy policies in the VR domain must tackle the following research challenges.

- **There is currently no unified criteria for vetting privacy policies, due to the decentralized and heterogeneous nature of VR platforms.** For a given type of VR device (e.g., Meta Quest 2), VR apps can be published on either its native platform (e.g., Meta) or compatible third-party platforms (e.g., SideQuest). However, the diverse regulatory requirements across heterogeneous app platforms (e.g., Meta requires a privacy policy for published apps while SideQuest does not) lead to varying quality in privacy policies, further complicating the setting of appropriate vetting criteria.
- **Existing vetting tools experience significant performance drops in VR domain.** Natural language processing (NLP) driven privacy tools (e.g., PolicyLint [5], PolicyGraph [23]) leverage named entity recognition (NER) and terminologization to process privacy policy sentences, but their performance is fragile when facing domain changes (e.g., shifting from mobile phones to VR). The emergence of many new data objects in VR (e.g., avatar, eye tracking) has negatively impacted the performance of these tools. For example, the NER of the latest privacy policy analysis tool, PoliGraph [23], only achieves an 87% recall rate for VR-domain privacy policies sentences, a stark contrast to its 95% recall rate on general-purpose privacy policies sentences. Current efforts to address domain changes either rely on purely manual checks (e.g., in smart home domain [41]) or lack comprehensive coverage of VR-specific terms (e.g., OVRSeen [71] only covers 100 privacy policies in VR).

VPVET. To address the above challenges, we develop VPVET, a novel system for vetting privacy policies for VR apps. To establish the vetting criteria, we survey the current VR platform market and select 10 mainstream platforms as our research focus. By crawling and collecting more than 11.9k apps' information, we construct the first and currently largest VR privacy policy dataset named VRPP. Then, based on the case studies from five typical VR apps' privacy policies and three representative privacy laws, we summarize five criteria for vetting VR privacy policies. Specifically, these criteria includes *availability* and structural *completeness* of a privacy policy, as well as their CUS statements' *granularity*, *minimization* requirements and *consistency* with actual behaviors.

To enhance domain-shift vetting performance, VPVET first automatically synthesizes a VR-domain policy sentences dataset. This includes 1.3k CUS sentences embedded with 267 unique VR domain phrases and an additional 14k non-CUS sentences. Utilizing this synthetic dataset, VPVET fine-tunes PrivBERT [69], a privacy

policy language model pre-trained on ~1 millions general privacy policies, resulting in a recall rate increase from 87% to 98.2% for VR-domain CUS sentences compared to the latest tool PoliGraph [23]. Additionally, VPVET introduces a semantic similarity-based clustering method that expands the terminologization coverage in the VR field by adding 84 more data object terms and covering an extra 5.8k phrases compared to OVRSeen [71]. Based on the larger number of data object phrases and more complicated terminologization, VPVET defines novel metrics (i.e., the lower bound and upper bound of privacy policy's claimed CUS) to assess the granularity of VR privacy policies and fairly analyze their minimization requirements and consistency with actual behaviors.

We leverage VPVET to measure privacy policies in VRPP and obtain following main findings (details are presented in Section 5). (1) *Inadequate availability and missing components*: platforms like Viveport and SideQuest have an availability rate even less than 1%. Furthermore, 65.3% privacy policies lack essential components, especially in relation to children's privacy statements. (2) *Coarse-grained CUS sentences*: 14.7% data objects (including sensitive data like health information) are not well specified. The disclosure of third-party collectors is even worse, with 93.5% of them not specified. (3) *Tendency of overbroad collecting information*: 85.1% VR apps present overbroad collections (similar to trends observed in Android apps [84]). In particular, big companies like Qantas and Emirates assign legacy privacy policies to their VR experience apps without any specification of VR traits. (4) *Discrepancy between policy and code practices*: 78.0% of apps we tested in VRPP show inconsistency between privacy policies and their actual code behaviors. What exacerbates this issue is that platforms like Meta have relaxed the inspection requirements for consistency, thus magnifying privacy breaches.

Contributions. We make the following contributions:

- We design VPVET to analyze privacy policies in the VR ecosystem, which covers the vetting criteria of availability, and structural completeness of the privacy policy, as well as granularity, minimization, and consistency of the CUS statements. Especially, the methodologies proposed by VPVET to handle the domain shifts in VR can be easily deployed for vetting privacy policies in other domains.
- We construct a large-scale dataset named VRPP, containing 11,923 distinct VR apps' meta-info from 10 mainstream VR platforms, as well as 3,521 valid privacy policies and 1,096 VR apps' package files. This dataset will be available to the public to facilitate further research.
- Using VPVET, we conduct the first large-scale measurement of privacy policies in VRPP. Our findings reveal significant mis-handling and disregard for privacy in the current VR ecosystem, as well as the potential reasons behind these phenomena.

Ethical consideration. Our analysis only relies on the meta-info and privacy policies of VR apps, which are publicly available (i.e., from VR platforms and homepages of VR apps). The dataset collection procedure are under the approval of the institutional review board (IRB) of our institutions. All discovered privacy issues are reported to the corresponding VR app platforms. We open-source VPVET and the findings to the public on <https://github.com/zyan-zhan/VPVET>




Figure 1: Sensors embedded in Meta Quest Pro (Left) and accessories emerging in the VR consumer market (Right).

//github.com/kalamoo/PPAudit to help the VR community better assess privacy policies.

2 BACKGROUND

2.1 VR Devices

Virtual reality is a cutting-edge technology that provides users with a fully immersive experience, attracting major tech companies to invest in developing their own brand of VR headsets. Take one of the most advanced standalone VR devices, Meta Quest Pro [44], as an example, it features 16 cameras, 29 infrared LEDs, 3 IMUs, and many other sensors [79] which are shown in Figure 1. Currently, there is a growing trend of integrating physiological and environmental monitoring accessories into VR devices [19]. Moreover, new accessories such as treadmills, haptic gloves, and brain interfaces are emerging with compatible consumer products already being introduced to the market.

According to their work mode, VR devices can be classified into four types: (1) *Optical Lens VR* consists of two convex lenses and several pieces of cardboards. (2) *Smartphone VR* works by connecting to a mobile phone and acts as display screens as well as I/O devices. (3) *PCVR* works much like Smartphone VR, but they are connected to a PC. (4) *Standalone VR* devices can work independently, allowing users to play in any area without potentially dangerous cables. Some Standalone VR devices are compatible with PCVR or Smartphone VR mode. More details can be found on Section 4.1.

2.2 Privacy Policy in VR

Requirements and expectations about privacy policy. A privacy policy is a statement that outlines how developers collect, use, share, and manage their users' data. The growth of privacy laws in recent years necessitates the privacy policy to function both as a notifications and a legal agreement between users and developers. According to [30], 182 jurisdictions have enacted over 1,043 specialized privacy laws, some well-known ones including PIPL in China [59], GDPR in Europe [27], CCPA [20] and California Privacy Rights Act (CPRA) [22] in California, Personal Information Protection and Electronic Documents Act (PIPEDA) in Canada [58], General Data Protection Law (LGPD) in Brazil [39] and Act on the Protection of Personal Information (APPI) in Japan [7]. Developers are responsible for complying with corresponding regulations by providing a publicly available privacy policy to their users and

properly informing users of data collection as well as their privacy rights. A widely adopted category method [4, 32, 64, 70, 78, 80] summarizes 7 components that a complete privacy policy is expected to include (See Table 1). Note that we have merged the 1st and 3rd party data collection categories into *Data CUS* while excluding *Do Not Track (DNT)* and *Others* categories.

Early stage of privacy policy in the VR ecosystem. It is reported that 91% users agree to the privacy policy by clicking on the checkbox without necessarily reading it [24, 43]. This phenomenon will have a more negative impact on VR scenarios: firstly, VR apps harvest more information about the user than existing mobile apps; secondly, some mainstream VR platforms pay little attention to vetting privacy policies and several platforms even do not require app developers to provide privacy policies when publishing apps. **Thus, it is desirable to design an effective and holistic tool to audit the privacy policies in the whole VR ecosystem.**

3 MOTIVATION AND VETTING CRITERIA

The goal of this study is vetting the privacy policies in VR ecosystem. To justify the rationale behind our vetting method, we begin by presenting a hypothetical yet realistic scenario that users may encounter given current VR technologies¹. Following this, we will summarize five vetting criteria for VPVET based on analyses of typical VR app privacy policy examples.

Motivation scenario: *Riley participated in a virtual reality maze and solved puzzles by physically walking around her living room. She was able to interact with her friends using gestures such as nodding, high-fiving, and making eye contact. Unbeknownst to Riley, her 20-minute VR game session captured 2 million data points of her body movements, which were sold to an insurance company. As a result, the company denied Riley's life insurance policy due to her behavioral patterns resembling early dementia. Her sister was also rejected for insurance policies as dementia tends to run in the family...*

Currently, there are no standards or regulations governing how VR data should be collected/used/shared. If Riley encounters the aforementioned situation where her data has been mishandled, she may first choose to read through the VR app's privacy policy for insights into what happened. However, Riley may face one of the following situations as shown in Figure 2, which inspired us to propose criteria for vetting VR apps.

Situation 1: She was unable to locate the privacy policy on either the app's homepage or the VR platform, like the case of the app SoundSphere published in the platform MicrosoftMR lacking a publicly available privacy policy.

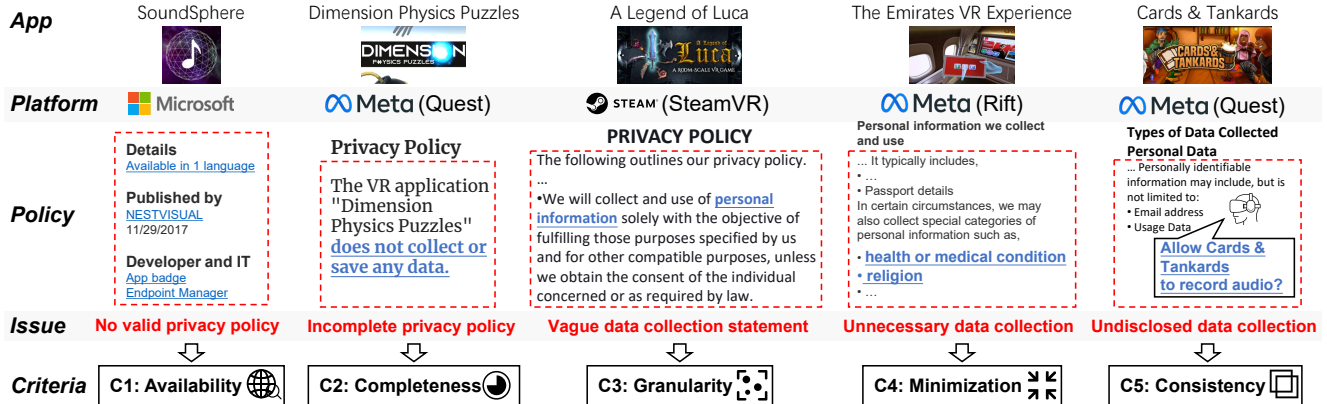
As such, our first vetting criterion is **whether the VR app provides an easily accessible privacy policy for users (C1: Availability)**.

Situation 2: She found the privacy policy, but it provides insufficient context for her to understand her privacy rights, like the case of the app Dimension Physics Puzzles published in Meta Quest.

¹This story originated from 2018 XR Privacy Summit [37].

Table 1: Explanation and examples of components in the privacy policy. Examples are extracted from the privacy policy of VRChat, the most popular social VR app attracting more than 19 million active users [21, 77].

Component	Explanation	Example
Data CUS	data be collected/used/shared	<i>we may collect or you may provide Personal Information about you</i>
Data Retention	for how long user data is stored	<i>when determining the specific retention period, we consider various factors</i>
Data Security	how user data is protected	<i>we use technical safeguards to improve the integrity and security of Personal Information</i>
User Choice	options available to users	<i>you may opt out from receiving commercial email by sending your request to us by email</i>
User Rights	users access/edit/delete their info	<i>you may submit a verifiable request that we delete Personal Information about you</i>
Policy Change	how users be informed about changes	<i>if we modify this Policy, we will make it available through the Platform</i>
Specific Audiences	pertain to specific groups	<i>we do not knowingly collect Personal Information from children</i>

**Figure 2: Motivation of VPVET: the current situation of VR app privacy policies and corresponding vetting criteria.**

Therefore, our second vetting criterion is **whether the VR app's privacy policy contains the necessary structural components (C2: Completeness)**.

Situation 3: She discovered that the privacy policy contains unclear statements regarding data collection (e.g., only stating personal information without further refinement in the privacy policy of A Legend of Luca, the best HTC Vive Games of 2016), raising her concerns about the specific type of personal information that will be collected or shared.

Therefore, our third vetting criterion is **whether the VR privacy policy provides detailed and specific statements about data collection in a clear and precise manner (C3: Granularity)**.

Situation 4: She found that the privacy policy claims to collect various types of data (e.g., medical conditions and religious info claimed in the privacy policy of The Emirates VR Experience published in the platform Rift) that do not seem necessary for the app's functionalities (i.e., VR experiences of a flight).

Therefore, our fourth vetting criterion is **whether the VR privacy policy only claims to collect the minimum amount of data required to support its functionalities (C4: Minimization)**.

Situation 5: She remembered granting several sensitive permissions to the VR app (e.g., microphone permission in Cards & Tankards, the most popular VR card game), but discovered that the privacy policy does not mention anything about it.

Therefore, our final vetting criterion is **whether the VR privacy policy aligns with its actual behaviors (C5: Consistency)**.

To summarize, there exists a significant disparity between the VR device's capacity to collect/use/share users' data and the insufficient transparency in informing users about it. With this motivation, we propose VPVET, the first automatic privacy policy vetting system for VR apps, to assess privacy policies based on the five aforementioned criteria.

Justification between vetting criteria and legal requirements. Some legal policies related to our vetting criteria are listed in Table 10 in Appendix B. It should be noted that failing to satisfy certain criteria does not necessarily mean that a privacy policy violates the corresponding legal articles listed in this table. Legal judgment is complex and requires consideration of several factors, making it beyond the scope of VPVET to provide legal advice or conclusions. The 5 vetting criteria and privacy reports generated by VPVET are intended only as recommendations and references to assist users in making informed choices and judgments.

4 VPVET SYSTEM

The workflow of VPVET is shown in Figure 3. We first provide a brief introduction to our data collection process in the VR app ecosystem and the availability vetting process in Section 4.1. Then, we demonstrate how to vet structural completeness in Section 4.2, followed by how to overcome the domain-shift challenges and accurately extract CUS tuples in Section 4.3. Finally, in Section 4.4, we show how VPVET defines the granularity metric for privacy policies in VR and vets the minimization and consistency based on granularity analysis.

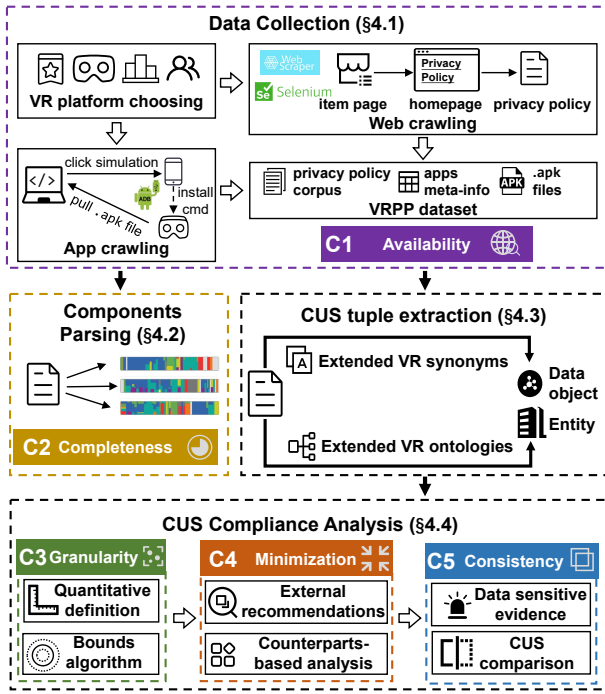


Figure 3: System overview of VPVET.

4.1 Data Collection (Vetting C1)

To vet privacy policies, VPVET first needs to collect a large-scale dataset that can reveal the true situation in the current VR app ecosystem. We select suitable VR platforms and then use web crawling techniques to retrieve their information and privacy policies. Additionally, to conduct the consistency analysis, we select a subset of these platforms (i.e., Standalone VR) and crawl their package files. The results of this module form the VRPP dataset, which is utilized in subsequent modules of VPVET.

VR platform selection. As depicted in Figure 4, VR platforms exhibit heterogeneous trends. Typically, each VR device is associated with a native VR content platform, alongside several third-party platforms available. For example, a Quest 2 owner can not only download applications from Meta Quest store, but also from third-party stores such as Sidequest or App Lab. Hence we start by searching for top-selling VR devices and include their native VR platforms as candidates. Next, we aggregate these candidates with compatible third-party platforms. Finally, we exclude Optics Lens VR and its corresponding platform, e.g., Google Play for the Google cardboard; we also exclude platforms that do not have public websites (e.g., Pico and Huawei VR, whose content is only accessible to VR device owners). After de-duplication, this process yields a total of 10 popular VR content platforms.

Extracting VR app meta-info and privacy policies. We utilize WebScraper [63] and Selenium [65] to crawl the meta-info (e.g., app name, app description, etc.) and privacy policy link of each app listed on its item page (i.e., the VR app’s info-page in certain platform). In cases where platforms do not provide direct links to privacy policies, we follow the homepage link (if available) of the




Figure 4: Mainstream VR devices and platforms.

app and extract clickable link objects whose text contains *privacy* as their privacy policy link. Finally, we download the HTML document for each privacy policy link using urllib [73] and convert them into plaintexts using HtmlToPlaintext [12].

Downloading standalone VR apps. Since only Standalone VR devices have their VR-specific apps while PCVR devices run desktop applications, we restrict the scope of app downloads to three Standalone VR platforms: the Official Quest store, its third-party app platform SideQuest, and the semi-official App Lab. For SideQuest apps, we use GetSidequestURL [53] to download the app’s package file through its source URL. For Meta Quest and App Lab platforms, since the app can only be downloaded through the official Quest app store on the VR headset, we automate this downloading process by simultaneously controlling a rooted Android mobile phone (with Meta Quest app installed and logging into our account) and the paired Quest device with a script running on PC.

In order to responsibly collect data and avoid impacting the VR platform server, we select a relatively slow speed and frequency of our automatic crawling method. The interval between each click for the web browser is longer than 10 seconds; the interval between each download command for the app’s package file is longer than 2 minutes; and the execution time of the command depends on the size of the app. Throughout our data collection procedure, we did not receive any complaints or warnings from these platforms.

Vetting availability (C1): we consider a privacy policy to be easily accessible and therefore *available* if it is provided either (1) on the item page of the VR platform, or (2) on the homepage of the VR app within two-hop. For a specific VR app platform, we define its availability as the ratio of published apps that have available privacy policies. The detailed vetting result are described in **FINDINGS 1** and **2** of Section 5.2.

4.2 Components Parsing (Vetting C2)

To vet structural completeness of a privacy policy, VPVET needs to parse the privacy policy into different components. Specifically, VPVET labels each sentence in the privacy policy according to its component category (as listed in Table 1) in the privacy policy. Due to the complexity of legal document text, a sentence may contain




Figure 5: CUS tuple extraction pipeline.

more than one type of statement. For instance, in the sentence “we may collect your IP address, if you do not want us to do so, you can opt-out”, both *Data CUS* and *User Choice* components are presented. Therefore, we model this classification problem as a multi-label task.

For this task, we deploy an end-to-end component parser consisting of a privacy policy language model called PrivBERT [69] (see Section 4.3.1 for more details) and a subsequent multi-label classification model. We train our parser on the OPP-115 dataset [78] using a 6:2:2 split for training, validation, and testing, employing the same hyper-parameters as PrivBERT. The components parser achieves an F1 score of 81.3% on the OPP-115 test set. To evaluate its performance in our VRPP dataset, we randomly select 200 sentences and manually evaluate classification results. The parser obtained a macro recall rate of 91.0%, demonstrating its efficiency. Finally, we deploy this parser to analyze all privacy policies in the VRPP dataset.

Vetting completeness (C2): A privacy policy is considered as *structurally complete* if it includes all 7 components (listed in Table 1), and **FINDING 3** of Section 5.3 shows the vetting result.

4.3 VR Domain-adapted CUS Extraction

To vet the compliance of CUS statements in the privacy policy, we first have to extract CUS tuples. A CUS tuple (e, d) comprises two components: an entity e and a data type d , indicating that the entity e has collected, used, or been shared with type d of data from users. The pipeline of extracting CUS tuples is shown in Figure 5, consisting of 4 main steps, i.e., CUS sentence identification, Named entity recognition (NER), CUS tuple extraction and Terminologization.

However, the domain changes from mobile apps’ privacy policies to VR apps’ privacy policies can negatively affect these steps’ performance, especially the recall rate of final VR specific terminologized CUS tuples. For example, for the CUS sentence “We may collect your ip address”, the CUS tuples within it can be extracted and terminologized by privacy policy analysis tools like OVRSeen [71] (which is based on PolicyLint [5]) and PoliGraph [23]. However, if we changes the CUS sentence to “We may collect your eye tracking data”, the corresponding CUS tuples although can be extracted but cannot be successfully terminologized. This is because even the VR oriented privacy policy tool OVRSeen doesn’t include or record *eye tracking data* in its synonyms files and ontologies². If we further complicate the CUS sentence with semantic structure like “There

may also be opportunities for you to grant permission for use of other of your eye tracking information”, then the latest general purpose privacy policy tool PoliGraph cannot recognize the data type using their NER.

Therefore, **VPVET** proposes a collective of techniques to enhance the VR-domain adaptability of this CUS tuple extraction pipeline. These includes: (1) constructing a synthetic dataset to boost the performance of CUS sentence identification and NER, (2) extending the share and collect (SoC) verbs for better CUS tuple extraction and (3) proposing a semantic similarity-based clustering method to efficiently enlarge the coverage of existing VR synonyms and ontologies.

4.3.1 Synthetic Dataset-booster CUS Sentence Identification and NER Models. VPVET first identifies those sentences that claim to collect/use/share users’ data (*CUS sentences*) from the privacy policy text and then label data object and entity tokens within each CUS sentence. Specifically, VPVET utilizes a synthetic dataset combined with the OVRSeen [71] synonyms file (with 267 VR domain phrases of 41 different types extracted from 100 VR privacy policies) and PI-Extract [16] dataset (with fine-grained manually annotated labels of data objects from 30 privacy policies) to train these two models. We retained sentences that do not contain any labels of data objects (*non-CUS sentences*) and considered others as *candidate CS sentences*. We then inserted VR domain phrases into the labeled position within these candidate CS sentences. In total, we obtained 14k non-CUS sentences and 1.3k CUS sentences embedded with VR domain terms. We train the models on the synthetic dataset using a 6:2:2 split for training, validation, and testing. As a result, the CUS sentence identification model achieves a F1 score of 82.0%, while the NER model achieved a F1 score of 86.5%. For comparison, the latest domain-shift NER for smart home privacy policies [41] achieves an F1 score of 75.75% with the assistance of a manually curated dataset of 284 privacy policies on that domain.

We attribute this performance improvement to the following two reasons. First is the **privacy policy language model** we use, i.e., PrivBERT that is pre-trained on millions of privacy policies which enables it to better capture features of privacy policies. Masked Language Modeling (MLM) task on VRPP-Corpus results indicate that PrivBERT has a better perplexity (8.59) compared to a general-purpose language model like distilroberta-base (10.86) used in PoliGraph [23]. Second is the **high-quality synthetic dataset**, which leverages VR-domain knowledge and enables the model to better recognize VR-domain entities. Additionally, the CUS sentences it utilizes is constructed from real-world privacy policies which often have more complex syntactic structures than simple statements

²The ontology is a directed tree, which serves to represent the subsumptive relationships between terms, where a link from term A to term B indicates that A is a broader term (hypernym) that subsumes B . More details can be found on Section 4.3.3.

Table 2: Extended SoC verbs used by VPVET.

Type	Word
Sharing	disclose, distribute, exchange, give, provide, rent, report, sell, send, share, trade, transfer, transmit, pass, express, supply, display, deliver, release, publish, lease, download, reveal, tell, view, show, hold, swap, forward
Collection	access, check, collect, gather, know, obtain, receive, save, store, use, get, perform, analyze, process, log, keep, add, record, combine, retain, recognize, track, remember, relate, create, ask, conduct, monitor, request, link, associate, solicit, read, preserve, contain

like “*We will collect your <data>*”. This helps the model recognize unseen entities as long as they appear in the proper syntactic slot in these CUS sentences. To evaluate this, we selected 47 VR-domain data objects (e.g., *body measure* and *arm length*) and 42 general data objects (e.g., *email address* and *age*), and insert them into the CUS sentences to test the latest general-purpose NER in PoliGraph [23] and our model. The results demonstrate that PoliGraph-NER can detect an average of 95.0% general data objects and 87.0% VR-domain data objects, while our models can achieve a detection rate of 96.5% for general cases and 98.2% for VR-domain cases. Therefore, this method avoids manually annotating hundreds of VR domain privacy policies while achieving great performance in handling VR-specific phrases.

4.3.2 CUS Tuples Extraction with Extended SoC Verbs. Given that a sentence may contain multiple data objects and entity tokens, we leverage the data and entity dependency (DED) tree from PolicyLint [5] to establish a mapping between the data object and their corresponding entities based on their grammatical relationships within the context. We also expand the list of sharing or collecting (SoC) verbs from 23 words to 64 words (see Table 2) to improve the recall rate of the DED tree. These improvements have resulted in the discovery of 54.5% (125,909/231,106) newly identified CUS tuples from VRPP-Corpus.

4.3.3 Terminologization. It is flexible for privacy policies to employ different phrasings when referring to the same data object or entity. For instance, phrases such as *record of your voice instruction* and *voice clip* both stand for data object term *audio*. Terminologizing these semantically similar phrases (i.e., *synonyms*) can help simplify subsequent analysis on CUS compliance. Existing works either uses manually curated lists of synonyms (PolicyLint [5] and OVRSeen [71]) or uses patterns-based method (PoliGraph [23]) to terminologize phrases. However, these methods either require massive human labor on checking every extracted phrases or will miss a significant number of phrases that do not comply with pre-defined patterns. For instance, when we utilize OVRSeen [71] (PoliGraph [23]) to terminologize extracted CUS tuples from the VRPP-Corpus, 7,069 (8,419) data object phrases, which covers 42,514 (49,254) CUS tuples, are reported as un-terminologized. Therefore, we propose an approach based on the insight that *synonyms have similar semantics and will therefore be clustered in the embedding space*. VPVET first utilizes a BERT-based sentence embedding model to map all phrases to the semantic embedding space. During this phase, any phrases that are within a threshold distance (0.8 in our study, i.e., the median similarity of the OVRSeen synonyms file) are added

Table 3: Comparison of OVRSeen and VPVET ontologies and synonyms.

Platform	OVRSeen	VPVET	New-in VPVET
Data ontology	63	107	84
#Data synonyms	2009	8042	5861
Entity Ontology	64	117	60
#Entity synonyms	894	1663	969

to synonym lists. For remaining un-terminologized phrases, we iteratively spot new clusters in embedding space and determine whether and where they can be included in the VR data ontology. Additionally, we employ keyword matching method to handle entity phrases that lack semantic meaning (such as the names of apps, developers, or domains). More details be found on Appendix A. As a result, we obtain an extended VR data (entity) ontology with 107 (117) nodes along with synonym lists containing 8,042 (1,663) distinct phrases (See Figure 11 of Appendix A and summary of changes in Table 3).

4.4 CUS Compliance Analysis

4.4.1 Granularity Analysis (Vetting C3). In some cases, privacy policies provide vague statements regarding data collection, rather than specifying the exact types of data that being collected, used or shared. For instance, compared to a CUS sentence saying it will “*collect your health information*”, one stating it will “*collect your health information such as your workout data*” provides users with a clearer understanding of what data is being collected. However, to the best of our knowledge, no existing metrics can quantitatively measure the granularity aspect of CUS statements in a privacy policy. To address this, VPVET introduces two metrics: *CUS tuple granularity*, which measures the granularity of the entity and the data type term in a give CUS tuple; and *privacy policy granularity*, which measures the granularity of a privacy policy by the set of data types they claim to collect. Their definitions are given below. **CUS tuple granularity (CTG).** As illustrated in Figure 6, for every node v (data object or entity) in the VR ontology O , the closer to the leaf node, the more fine-grained it is. Hence, we define its granularity as the longest distance from it to any leaf node:

$$CTG_O(v) \stackrel{\text{def}}{=} \max_{s \in S} \left(\max_{p \in Paths_O(v,s)} (Len_O(p)) \right), \quad (1)$$

where S is the set of all leaves in O , $Paths_O(v,s)$ is the set of all simple paths from node v to leaf s in O , and $Len_O(p)$ is the length of path p (defined as the number of nodes in this path) in O . The range of CTG of a node in VR data or entity ontology (See Figure 11) is from 1 (the leaf) to 5 (the root). The smaller the CTG, the more fine-grained the term. For example, in Figure 6, $CTG_O(workout) = 1$ while $CTG_O(health) = 2$. Note that, for a given CUS tuple, since it contains both an entity e and a data object d , the calculated CTG is a 2-tuple $\langle CTG_O(e), CTG_O(d) \rangle$.

Privacy policy granularity (PPG).

We define this based on the following two insights. *Firstly*, claimed CUS tuples in a privacy policy have two bounds: the lower bound (data types that are explicitly stated as being collected) and the




Figure 6: Granularity illustration. This is a snippet from the complete data ontology shown in Figure 11.

upper bound (potential data types that may be collected). For example, in the CUS sentence “we may collect your health information such as your workout data”, *workout data* is explicitly claimed to be collected while the other types of data like *blood sugar* will be potentially collected because they are the child of the *health* node in the ontology. Figure 6 illustrates the lower and upper bounds of this example. See Algorithm 1 and Algorithm 2 in Appendix for more details on how to calculate them. These bounds satisfy the following inequation:

$$\emptyset \sqsubseteq \text{LowerBound} \sqsubseteq \text{Claimed} \sqsubseteq \text{UpperBound} \sqsubseteq \mathbb{U}_T, \quad (2)$$

where \mathbb{U}_T is the set of all terms in the ontology. Secondly, if the gap between the lower and upper bounds is small, we consider an app’s privacy policy to be fine-grained, as there are no ambiguous interpretations for these statements. Therefore, we define the PPG of the given privacy policy as

$$\text{PPG} = \text{UpperBound} - \text{LowerBound}, \quad (3)$$

The range of PPG is from 0 to the number of nodes in the data ontology (i.e., 107 in this study). The smaller the PPG, the more fine-grained the privacy policy.

Vetting granularity (C3): We regard the CUS tuple granularity as coarse-grained if the CTG is equal or larger than 2. As for the privacy policy granularity, considering PPG is a relative value, we will not set a threshold to determine whether an entire privacy policy is fine-grained or not. Instead, we’ll provide PPG’s distribution in **FINDING 4** of Section 5.4 and PPG’s ranking percentiles in each app’s privacy policy report. We also provide the vetting results regarding granularity aggregated by different VR platforms.

4.4.2 Minimization Compliance Analysis (Vetting C4). Minimization compliance requires a privacy policy to only collect necessary data from users. However, there is no legal definition of *minimization* specifically for various apps. To address this issue, we adopted the counterpart-based method described in [84], which compares CUS between a target app and similar apps (referred to as *counterparts*). The insight behind this approach is that, if a group of apps provide similar functionalities, then they are expected to collect similar scope of data to support those functions. Under this assumption, any CUS data that is not collected by the majority of counterparts can be considered unnecessary, i.e., *overbroad*. Considering that privacy policies have different granularity levels when referring to a data type, we perform such counterparts-based comparisons based on their *LowerBound* of CUS statements.

To find proper counterparts for each target app, we propose a multi-sources counterpart searching that considers (1) direct recommendations within a single platform (such as SteamVR), (2) cross-platform recommendations from professional recommend

Table 4: VRPP-Corpus description.

Platform	# APP Info	# PP	$\frac{\#PP}{\#AppInfo}$
Sidequest	2,274	192	0.084
Viveport	2,919	281	0.096
PSVR	580	91	0.157
SteamVR	6,748	1,185	0.176
Microsoft	285	83	0.291
Gear	1,085	1,083	0.998
Go	1,118	1,116	0.998
Rift	1,368	1,366	0.999
Quest	387	387	1
App Lab	1,335	1,335	1
De-duplicated Summary	11,923	3,521	0.295

Table 5: VRPP-APK description.

Platform	# APP Info	# PP	# APK	# APK w/ PP
Sidequest	1,319	83	691	46
Quest	116	108	41	36
App Lab	981	929	364	349
De-duplicated Summary	2,416	928	1,096	286

websites (like steampeek), (3) app genres, and finally determine the top-k counterparts based on their (4) app descriptions.

Vetting minimization (C4): We consider a privacy policy to meet the minimization criterion if none of its claimed CUS tuples has overbroad data type. Vetting results regarding minimization aggregated by data types as well as aggregated by different VR platforms can be found in **FINDING 5** of Section 5.4.

4.4.3 Consistency Compliance Analysis (Vetting C5). The app is expected to disclose all collected data objects in its privacy policy to ensure consistency compliance. As we do not have access to the server database of the app, we examine its behavior based on its code by de-compiling source code from their apk files.

We focus on data-sensitive evidence in the source code. This evidence includes sensitive permissions required by the app and data-sensitive functions/methods/APIs/URIs. We manually constructed a mapping file that illustrates the mapping relation between data objects and sensitive permissions and APIs on VR apps based on [9]. We first updated this mapping to be compatible with Android 10 and Android 12, which most standalone VR devices are built upon. Third-party VR APIs such as Oculus, WaveVR, and Samsung are also included to cover more VR-domain data-sensitive behaviors. In total, we construct a mapping from 167 data-sensitive evidences to 28 data objects (15 of which are of VR traits). Finally, we compared the data-sensitive evidence in code with CUS tuples extracted from an app’s privacy policy to examine their consistency.

Vetting consistency (C5): We consider data objects that have evidence in the app’s code to be vaguely claimed if they are in *UpperBound* – *Claimed* set. If a data object is not even covered by the *UpperBound*, then we regard it as inconsistent. The results can be found in **FINDING 6** in Section 5.4.

5 VETTING RESULTS

5.1 Dataset Description

Our dataset VRPP consists of two parts: VRPP-Corpus and VRPP-APK. VRPP-Corpus is collected during September 2022 and October 2022, including meta-info of 17,299 VR apps from 10 mainstream VR platforms. Meta-info includes apps' common attributes, such as name, platform, item page link, homepage link (if any), privacy policy link, publisher details, genre description, and price. Considering that a VR app may be published on more than one platform, we aggregate (de-duplicate after merging) information for identical apps (by name), resulting in a total of 11,923 unique apps. After preprocessing the privacy policy links to extract plaintexts, we obtained 3,521 valid privacy policy texts for further analysis. The distribution of apps and privacy policies in different VR platforms is shown in Table 4.

Since VRPP-Corpus can only support vetting for criteria C1-C4, to vet consistency (C5), we additionally collect VR package files and construct VRPP-APK, during September 2023. During dataset construction, we focus on free apps from Meta Quest, Sidequest, and App Lab platforms. This is the subset of the aforementioned dataset with an app file in .apk format. In total, we collected meta-info from 2,416 apps and obtained 928 valid privacy policies as well as 1,096 apk files (See 5). Among them, there are in total 286 VR apps with both .apk file and valid privacy policy downloaded. It should be noted that we failed on approximately 40.0% of the apk download links when trying to get the valid apk files on Sidequest platform. Moreover, there is a failure rate of 66.7% for downloading apps' apk files from Quest headset for Quest and App Lab platforms. We will discuss this limitation on Section 6.

5.2 Vetting Results of Availability

FINDING 1: Several VR platforms, including major ones like PSVR and Steam VR, due to inadequate regulations, have poor availability (less than 0.3) of privacy policies.

As shown in Table 4, for totally 11,923 VR apps we crawled on VR platforms, only a small portion (i.e., 29.5%) of privacy policies were successfully found. Specifically, half of the mainstream platforms in the VRPP-Corpus, named Viveport, Microsoft, PSVR, SteamVR, and Sidequest have low availability rates. On the contrary, App Lab, Quest, Go, Rift, and Gear - all under Meta's supervision - have high availability ratios with values close or equal to 1. Some real-world examples of VR apps fail to provide privacy policies can be found on Table 11 of Appendix C.

Potential reasons of Finding 1. We investigated the reasons behind this phenomenon and discovered that whether developers provide a privacy policy for their VR app largely depends on the platform's requirements. Meta provides guidelines to developers [49], in which they require developers to "maintain a publicly available link to your privacy policy ... and ensure the link remains current and up to date". This explains why platforms like App Lab, Quest, Go, Rift that are regulated by Meta has availability rate large than 99.8%. Similar guidelines can also be found on Microsoft and Viveport, however, they are not mandatory. Microsoft designates a *Privacy Policy URL as Required* in their checklist but notes that it is *Sometimes not require* for publishing an app [51]. Viveport states that "If you (developer) have your own privacy policy, you may enter




Figure 7: Reused privacy policies distribution.

its URL." but does not make it compulsory [75]. For the Sidequest platform, which performs worst in the availability of the privacy policy, we attribute it to its third-party status. In their terms [68], they claim to be "a listing service only and have no responsibility for the content or accuracy, completeness, or lawfulness of the listings or the games".

FINDING 2: Privacy policy reuse is a common issue for VR apps, with 54.5% of vetted policies being shared among different apps, despite variations in their data collection practices.

While processing the VRPP-Corpus, we discovered that several VR apps may choose the same document (referred to as *reused policy document*) as their privacy policies even without modifying any content. Figure 7 shows the number of reused privacy policies and how often they are reused. It is observed for 3,521 apps providing privacy policies, 54.5% (1919/3521) of them show the reuse behavior, with 687 unique policy documents being reused a total of 1919 times. Note that this statistic represents the minimum number of instances where privacy policies are reused, as some policies may also be used by other apps (including non-VR apps) beyond the scope of our VRPP-Corpus. Two policies stood out, each with more than 15 reuses: Oculus's privacy policy document [48] (reused by 19 VR apps developed by Oculus) and Adobe's privacy policy document [3] (reused by 16 VR apps shared on Behance [11], a social media platform owned by Adobe that focuses on showcasing creative work). We also identify the privacy policy of 17 different VR apps all re-direct to Microsoft's privacy policy [52].

Privacy risks of reuse behavior of privacy policies.

A common pattern among these reused behaviors is that a certain developer publishes multiple VR apps and applies the same document as policy policies to all of them. In this scenario, the CUS claimed in the privacy policy typically is the union set of actual behaviors across all VR apps. As a result, despite the fact that two apps may have different data collection practices, users will see identical CUS in their reused privacy policy. For instance, *VZfit* (one of the top VR fitness apps with 611 ratings and rated 4/5) and *VZplay* (with 2.3k clicks and 16.2k views, rated 5/5 on Sidequest) are developed by VirZOOM Inc. and share the same privacy policy; however, they differ significantly in terms of their data collection practices, i.e., *VZfit* requires fine-grained location access and microphone permission from users, while *VZplay* does not request these permissions but accesses hand tracking data instead. Another example is Oculus, which not only functions as a platform but also

operates as a developer. Within our VRPP dataset, Oculus has released 19 official apps, all direct to the same privacy policy [48]. However, most of these apps display varying data collection behaviors. For example, *The World Beyond* (with 5.8k clicks and 19.5k views, rated 4.6/5 on Sidequest) collects microphone data and requests coarse-grained location access and hand-tracking data; whereas *Dear Angelica* (an interactive VR story developed by official Oculus Story Studio, ranked 9th among top-rated free VR apps with 281 ratings and an average rating of 4.35/5) does not have such practices.

From a commercial standpoint, it is cost-effective and legally sound for large companies or developers to maintain a single but comprehensive privacy policy that covers all possible data collection for their services and apps, regardless of the VR features or variations among different VR apps. However, this approach may lead to user confusion and breed mistrust.

5.3 Vetting Results of Completeness

FINDING 3: 41.7% VR privacy policies did not adequately inform users about their privacy rights and 65.9% of relevant VR apps failed to address children’s privacy concerns.

As shown in Table 6, for all vetted privacy policies, only 34.7% out of them provide all necessary 7 components described in Table 1. Over 81% of privacy policies contain statements about data CUS, but less than 66% of them provide other important components such as informing users’ rights (58.3%) to access, edit, and delete their personal data on the app’s server, statements about data retention (57.1%), and specific audiences (61.2%). Some real-world examples of VR apps fail to provide complete privacy policies can be found on Table 12 of Appendix C. In the below, we present two case studies about policies without any valid component and improper handling of children’s policy.

Case study: privacy policy with no valid component. Out of VRPP-Corpus, only 34.7% (1223/3521) of privacy policies provide ALL these components while surprisingly, 11.5% (406/3521) provide NONE of these components. Upon manual inspection of 50 randomly selected privacy policies that lacked all these components mentioned above, we summarize three types of situations: (1) fake-redirection, where clicking on the privacy policy link redirects to the same page instead of the privacy policy content page, e.g., from <https://sloppystudio.com/> to <https://sloppystudio.com/#> on *Car Parking Simulator VR* (rated as *Mostly Positive* in SteamVR); (2) ongoing privacy policy, where the page displays “*Under Construction*” with no additional information, e.g., *Chandrayaan VR*; and (3) dummy privacy policy, where only one sentence is displayed, such as the privacy policy [1] “*No user data is collected from the app manufacturer*” for *Las Vegas Helicopter Flight* and *Venice - Grand Canal*, both priced at 1.99 USD on Quest Store. While these privacy policies may maintain valid links, they do not provide any useful information.

Case study: children’s policy. Given the popularity of immersive gaming experiences and the widespread use of VR devices by young children [14, 61], it is crucially important for VR app developers to address privacy concerns related to children in order to comply with children’s data protection laws (e.g., COPPA [26]). We focused on apps related to education, kids (children), or family genres and

Table 6: Ratio of privacy policies with necessary components.

Component	Ratio(%)	Component	Ratio(%)
User Choice	65.6	Data Security	65.6
Data CUS	81.1	Policy Change	64.2
User Rights	58.3	Spec. Audience	61.2
Data Retention	57.1	All 7 Components	34.7




Figure 8: Distributions of CUS tuple granularity and privacy policy granularity.

checked for specific statements regarding children’s privacy. Out of 718 children-related apps we find, 65.9% (473/718) do not mention anything about children’s privacy in their policies. For instance, *Bogo* (a pet feeding game with 1326 ratings, ranked 8th among the top free VR apps with the highest number of ratings, and has an average rating of 4.2 out of 5), *Henry* (a storytelling app with 394 ratings and rated 4.0/5), and *Paper Birds* (an interactive story with 268 ratings and rated 4.4/5) are VR apps published on Quest and designed for users aged 3+. However, none of them contain specific statements about children’s privacy in their policies. Specifically, the policy of *Paper Birds* only states that they “*don’t collect any of your personal info at any time, ... have never received any legal or government demands for user information*”. A similar situation arises with two other representative VR apps, namely *Pets VR* (with 78.6k clicks and 264.4k views, rated 4.2/5 on Sidequest) and *Ultimate Fishing Simulator VR* (which has received 501 *Very Positive* reviews on SteamVR).

5.4 Vetting Results of Granularity

FINDING 4: Some VR privacy policies lack fine granularity in both data object disclosure and third-party specification.

For totally 25,895 CUS tuples extracted from privacy policies in VRPP-Corpus, we first calculate the PPG of each privacy policy. Then we categorize all these CUS tuples into two groups: 17,487 first-party CUS tuples (where the entity is *we* and the entity granularity equals 1) and 8,408 third-party CUS tuples (where the entity is explicitly or implicitly claimed as *third-party*). We then calculate CTG for each individual CUS tuple. The results are shown in Figure 8.

Granularity of privacy policy is coarse. Figure 8(a) displays the CDF of PPG for all valid privacy policies. According to this figure, the inclusion of unspecified data objects in a privacy policy significantly increases the upper bound of claimed data objects. On

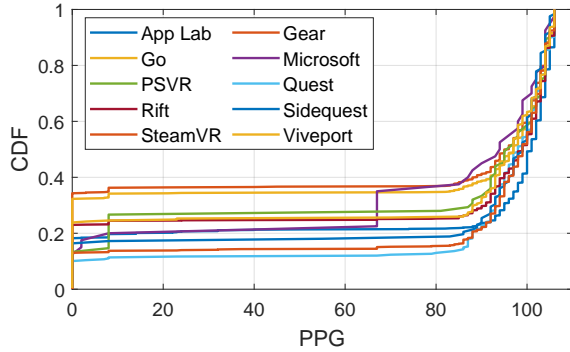


Figure 9: PPG over privacy policies on different VR platforms.

average, a privacy policy will include 75.4 data objects within the gap between its lower and upper bound. The median value is 99.0. Therefore, we can conclude that VR apps' privacy policies tend to use coarse-grained CUS statements rather than fine-grained ones.

First-party CUS provides coarse-grained disclosure of data objects. It is observed from Figure 8(b) that 8.5% (1,489/17,487) of first-party CUS tuples have a data object granularity higher than expectations (i.e., equal or larger than 2 in this study). Considering the definition of lower-bound, this indicates that these CUS are claimed *without* any additional clarification or introduction in the privacy policy. Here are several typical examples: 18 privacy policies (including Intel's privacy policy [36] for its VR app *Queerskins: Ark* in Viveport) fail to clarify the meaning of *Biometric data*. Additionally, 79 privacy policies (including the well-known meditation and sleep VR app *CalM*) do not provide a clear definition for *Health information*. Furthermore, 43 privacy policies, which include the popular game *Bloppy Tennis* (with 48k clicks and 94.5k views) and the VR social platform *Cheerio* (with 2.6k clicks and 15.7k views), do not specify what is meant by *Body Measurement*.

Granularity of third-party CUS is even worse. 27.6% (2,318/8,408) of third-party CUS tuples do not specify the exact type of data object they collect, and surprisingly 93.5% (7,861/8,408) of CUS tuples do not mention the exact company name of the third-party. The most common case is to refer to them as *third-party* (4106). In other cases, they may be categorized as specific types of third parties such as *ad network* (708) and *platform provider* (578). Among the explicitly mentioned third parties, Google (186), Google Analytics (77), Unity (102), and Facebook (92) are the most prominent. It is worth noting that only 4.3% (365/8,408) of these third-party CUS clearly identify both the company name of a third party and the specific types of data collected by that particular third party.

Granularity vetting results varies among different VR platforms. Figure 9 and Table 7 display the differences of PPG results over privacy policies published on different VR platforms. Although there is little difference among their PPG-median values (from 94 to 101), in terms of PPG-mean value, apps in platform Gear averagely provide the most fine-grained privacy policies (with PPG-mean value as 63.57) while apps in platform Quest provide the least fine-grained granularity (with PPG-mean value as 86.89).

Table 7: PPG mean and median of privacy policies on different VR platforms.

Platform	PPG-mean	PPG-median	Platform	PPG-mean	PPG-median
Gear	63.57	96	Rift	75.37	99
Go	65.81	97	App Lab	80.23	101
PSVR	73.79	96	Sidequest	81.57	99
Viveport	74.11	97	SteamVR	84.88	99
Microsoft	74.53	94	Quest	86.89	98

5.5 Vetting Results of Minimization

FINDING 5: 91.6% of data objects and 85.1% of VR apps exhibit at least one overbroad situation.

In total, we find counterparts for 2,033 VR apps, and the minimization analysis results are shown in Figure 10, which displays the identified overbroad data objects (i.e., surpassing their counterparts in CUS) along with their frequency and ratio. Out of all data objects in our ontology, we find that 91.6% (98/107) are involved in at least one case of overbroad CUS. Similarly, out of all comparable VR apps (i.e., apps that we successfully find their counterparts), we find 85.1% (1699/1997) had at least one overbroad data object in their CUS tuples. These findings demonstrate the prevalence of overbroad situations and suggest a need for stricter adherence to privacy policies' minimization principle.

As depicted in Figure 10, 94.9% (93/98) data objects have an overbroad ratio exceeding 0.5 (meaning a fifty-fifty chance of being overbroad when claimed in a VR app's privacy policy). The overbroad ratio of 65.3% (64/98) of the data objects even exceeds 90%, and there are 29 data objects that are consistently identified to be overbroad. It is worth noting that some VR-related data objects such as *gameplay* (362/79.7%), *audio information* (231/94.7%), *camera information* (109/96.5%), and *VR movement* (83/98.8%) have both a large number of overbroad CUS cases and high overbroad ratios.

Case study of Finding 5: All-encompassing privacy policy of large company used on VR app. During our evaluation of the overbroad situation in privacy policies, we discovered that some VR apps link to a parent company's or association's privacy policy (referred to as an *all-encompassing* privacy policy). Note that this phenomenon is different from what we refer to as *reused* privacy policies in **Finding 2**, which are shared among different VR apps and include the union set of these VR apps' CUS. The *all-encompassing* privacy policies, instead, include all CUS requirements for the parent company's business affairs, without making any modifications or adding specific statements that apply to the traits of the VR app. Consequently, these privacy policies end up claiming more data objects than what the VR app actually collects. As a result, our evaluation highlights these policies as failing to meet the minimization criterion.

Taken *Qantas VR* (published in SteamVR) and *The Emirates VR Experience* (published in Rift) as examples, they claim to collect highly sensitive information such as ethnicity, beliefs, and passport details - unnecessary and impractical for flight experience VR apps. It turns out that these two apps link their privacy policy to the policy of their respective parent companies (*Qantas* and *Emirates*), where large amounts of personal information are claimed to be collected. Similar situations occur with *Lusail Stadium VR Experience* (published in Go and Gear) linking to *Qatar 2022 FIFA World Cup's*




Figure 10: Distribution of overbroad data objects and their overbroad ratio.

Table 8: Minimization vetting results of privacy policies on different VR platforms.

Platform	#Comparable	#Overbroad	$\frac{\#Overbroad}{\#Comparable}$
PSVR	62	46	0.742
App Lab	243	183	0.753
Sidequest	20	16	0.800
Viveport	196	168	0.857
SteamVR	685	589	0.860
Rift	807	696	0.862
Go	516	456	0.884
Gear	481	426	0.886
Quest	252	230	0.913
Microsoft	29	28	0.966

privacy policy, claiming to collect user passport information; and BYU Virtual Campus (published in SteamVR) linking to Brigham Young University’s (the largest church university in the USA) privacy policy, claiming to collect user education, employment, belief, and mental health information. Other real-world examples can be found in Table 13 of Appendix C.

Minimization vetting results varies among different VR platforms. Table 8 displays the differences of minimization results of those privacy policies published on different VR platforms. Since the sizes and the sources for searching for counterparts are different for different VR platforms, the number of VR apps that are *comparable* for minimization vetting would be different. Among all the platforms, PSVR has the lowest overbroad ratio, indicating that VR apps on PSVR have the *most minimized* privacy policies. In contrast, MicrosoftMR has the highest overbroad ratio, with 28 out of 29 comparable VR apps deemed to be overbroad.

5.6 Vetting Results of Consistency

FINDING 6: Inconsistencies between actual code behavior and privacy policies are common in VR apps, particularly with regard to VR-related data objects.

This analysis is based on VRPP-APK. We first observe that the issue of poor accessibility of privacy policies, as mentioned in **Finding 1** about Sidequest, still persists. Out of all 1,096 apps, 810 (662 of which come from Sidequest) do not provide the required privacy policy. Furthermore, the APK files of their apps exhibit evidence of sensitive permissions or behaviors, as indicated in the first column of Table 9. This could potentially breach the relevant privacy law (if applicable) because there is a discrepancy between the privacy policy (NULL here) and the actual behavior of the app.

Then we analyze the inconsistency of those apps (a total of 286) with both privacy policy and apk file available and the results are shown in the right part of Table 9. We found that 85.3% (244/286) of these apps vaguely claim their code behavior, while 15.7% (45/286) do not disclose their code behavior in the privacy policy. From the perspective of per code behavior, the overall inconsistency ratio is 12.4% (147/1183), with the majority of them (74.8%, or 110/147) related to three common data objects: network information, geo-location, and device information. For camera and hand tracking data objects, although the total number of inconsistencies is not significant, the relative inconsistency ratio in these data objects is higher than average, with 13.7% and 13.5% respectively. It is worth noting that we once again confirm the coarse granularity of CUS in VR app privacy policies. We have observed that 68.4% CUS practices in VR app code are not explicitly mentioned in their corresponding privacy policy. Instead, they are covered by a vague CUS statement with coarse granularity. Surprisingly, over half of the data objects (14/27) listed in Table 9 show vagueness across all their occurrences in VR apps. Among these, 71.4% (10/14) are due to highly VR-related data objects such as *vr play area* and *pupil distance*.

Current platform pay insufficient attention to policy’s vagueness and inconsistency. Towards this phenomenon, we might find some relevant explanations on the official Meta website. Meta considers a privacy policy necessary in their VRC (Virtual Reality Check) and conducts a detailed Data Use Checkup (DUC) when

Table 9: Evaluation results of consistency analysis. The cumulative results are shown in the Total row. For the #App, they are further de-duplicated based on names.

Data object	w/o pp	w/ pp		
	# App	# App	# / % Vague	# / % Inconst.
network	782	286	221 / 77.3%	39 / 13.6%
geo location	732	262	149 / 56.9%	35 / 13.4%
device info	703	252	165 / 65.5%	36 / 14.3%
audio	412	118	67 / 56.8%	11 / 9.3%
camera	164	51	42 / 82.4%	7 / 13.7%
hand tracking	128	37	31 / 83.8%	5 / 13.5%
billing	71	50	35 / 70.0%	5 / 10.0%
account	58	27	19 / 70.4%	5 / 18.5%
usage info	45	20	6 / 30.0%	1 / 5.0%
vibrator	26	12	12 / 100.0%	–
ad id	26	12	10 / 83.3%	1 / 8.3%
vr play area	5	7	7 / 100.0%	–
infrared	2	7	7 / 100.0%	–
eye tracking	2	6	5 / 83.3%	–
contact	2	–	–	–
pupil distance	1	6	6 / 100.0%	–
face	1	5	4 / 80.0%	–
accelerometer	1	2	2 / 100.0%	–
body measure	1	2	2 / 100.0%	–
gyroscope	1	2	2 / 100.0%	–
error report	1	2	2 / 100.0%	–
vr movement	1	2	2 / 100.0%	–
environment	1	4	4 / 100.0%	–
battery	1	1	1 / 100.0%	–
heart rate	1	1	1 / 100.0%	–
fingerprint	1	1	1 / 100.0%	–
biometric	1	1	1 / 100.0%	–
Total	810	286	809 / 68.4%	147 / 12.4%

reviewing an app. However, we discover that it informs developers “It does not need to provide the explicit data (e.g. Username, Profile picture, leaderboards) and can reference it in abstract” in their second and third criteria for the privacy policy [45, 47]. This suggests that Meta may encourage developers to use coarse-grained data collection statements even if both Meta and the developers are aware of the specific data types collected from users.

6 DISCUSSIONS

Impact of privacy policy correctness on vetting results. Except for the consistency and availability criteria, the other three criteria (completeness, granularity and minimization) are designed to vet the *quality* rather than the *correctness* of VR app’s privacy policy. For instance, consider a maliciously curated privacy policy. Assume this privacy policy: 1) includes a *Data Security* component but its VR app does not implement the claimed data protection measures, and 2) specifically states that it collects very limited types of user data, while in reality, the VR app attempts to collect as much user data as possible. As a result, this incorrect privacy policy would pass all the completeness, granularity and minimization vetting criteria and be considered as a *high-quality* privacy policy, even though the VR app’s actual behavior does not align with it. Although our consistency criteria partially address this issue, the challenge of

verifying the correctness of all components of a privacy policy remains an open problem.

Limited collection of privacy policy and package file. For privacy policies, we only considered easily accessible ones in Section 4.1, thus may ignore the cases where the privacy policy was not readily displayed on the homepage but could be accessed by, for example, adding a suffix like */privacy-policy* to the URL. We may also miss those privacy policies that use keywords (e.g., “statement”, “notice”, “legal”, “terms”, “agreement”, “disclaimer”, and “policy”) other than “privacy” on their homepage. Out of randomly selected home pages of 1k VR apps, we identified 14 privacy policies belong to this situation. For package files, we encountered significant network failures when attempting to download them. A simple retry can increase the number of successfully downloaded apps from 350 to 691. Therefore, we repeated our method three times and obtained the results in Section 5.1. We leave enhancing the data collection method’s robustness for future work.

Context-aware CUS tuple extraction. Our CUS tuple model only considers the entity and data object of a CUS, ignoring other contexts. For instance, it is common to put a CUS sentence like “when you *<condition>*, *<entity>* may collect your *<data obj>* to *<purpose>*” where extended CUS tuple can provide context integrity (CI) [8, 67] for privacy handling. This extended CUS tuple can help analyze privacy policy in a more holistic approach. However, automatically extracting context with high accuracy remains an open problem, and we will leave it as future work.

Incomplete coverage of VR apps’ behaviors. Firstly, the analysis is limited to apps that satisfy both StandaloneVR and Free conditions. Secondly, there are limitations in the coverage of app behaviors as sensitive permissions and APIs may include deprecated or unreachable dead code, resulting in false positives. Additionally, data objects covered by permissions or APIs are incomplete compared to the full data object ontology. Certain sensitive data categories such as mental health and preferences cannot be identified merely through code, but rather require a combination of other data types for inference. It may need to explore the intricate relationships between data types in VR scenarios and develop more precise methods for modeling VR app behaviors. We believe that an automatic tool capable of analyzing VR apps both statically and dynamically is necessary to cover as many VR data collection behaviors as possible. A promising direction would be developing a VR device simulator that can run VR APPs and inject inputs that simulate human interactions.

7 RELATED WORKS

Privacy issues in VR. The privacy concerns of VR apps stem from their immersive experience and the ability to collect sensitive data [25, 31, 38]. Various sensor data, such as motion, eye tracking, hand tracking, indoor tracking, facial capture, and body measurement [15, 18], can be used to infer users’ health or profile information with high accuracy [10, 54]. A semi-structured interview with 20 VR users and developers [2] explore the privacy perception in the VR community and co-design a set of ethics regarding VR content, which could potentially serve as industry-wide standards.

CUS extraction of privacy policies. There are two main technical approaches to extracting CUS statements from privacy policies.

First is the rule-based extraction (RBE) solution proposed in PoliLint [5] (also adopted by PoliCheck [6] and OVRSeen [71]) and PoliGraph [23]. PoliLint utilizes data-entity-dependency (DED) patterns to identify and extract CUS tuples from sentences, but it has a relatively low recall rate (<30%) due to limited pre-defined patterns. PoliGraph also extracts CUS tuples based on several rules. To improve the recall rate, it has to decompose this task into 5 relation-annotator tasks, including collection, subsumption, and coreference relations, and combines the results of these tasks on a graph structure. Second is the end-to-end extraction, as proposed by PI-Extract [16]. However, the CUS tuple extracted through this method lacks accurate entity information and only separates them into first-party or third-party collections.

Privacy policies in other domain. In addition to websites and mobile phone apps, privacy policies are required in many other domains. Perez *et al.* [57] and Yu *et al.* [83] have examined the privacy policies of IoT devices and their consistency with actual app behaviors. Manandhar *et al.* [41] did a thorough survey of privacy policies in the Smart Home domain, identifying 17 findings that impact millions of users. Bui *et al.* [17] proposed ExtPrivA, an automatic tool that detects inconsistencies between browser extensions' data collection and their privacy policies. Other studies focus on specific categories of applications such as menstrual apps [66], or money services [13]. Trimananda *et al.* developed OVRSeen [71] to audit network traffic and privacy policies in Oculus VR. However, in addition to the limitation of CUS extraction discussed above, OVRSeen covers only 150 VR apps and ignores other criteria for vetting a policy except for consistency.

Privacy policy generator. Drafting a privacy policy can be complex, which is why there are tools available to generate policies for Android apps [81, 82], iOS apps [85], and smart home apps [40]. Google announced Checks [29] to help developers draft and rectify their privacy policies according to relevant laws. However, these tools face challenges in extracting reasons for data collection practices from source code. Besides, some components like user rights cannot be inferred from app source code alone.

8 CONCLUSION

This study proposes VPVET, the first large-scale comprehensive privacy policy vetting system for VR apps. VPVET collects 11,923 different VR apps' meta-info and analyzes 3,521 valid privacy policies (from 10 mainstream VR platforms) as well as 1,096 VR apps' package files (from 3 VR platforms) based on the following 5 criteria: availability, completeness, granularity, minimization, and consistency. Our findings expose the significant mishandling and disregard for privacy within the current VR ecosystem. To increase awareness of the VR community and facilitate future analysis of VR privacy policies, we open-source VPVET system as well as our research findings on <https://github.com/kalamoo/PPAudit>.

ACKNOWLEDGEMENTS

The authors from Shanghai Jiao Tong University were partially supported by the National Natural Science Foundation of China (No. 62325207, 62132013, 62302298, 62332013), Young Elite Scientists Sponsorship Program by CAST (YESS20230589), and Startup Fund for Young Faculty at SJTU (23X010502192). The authors from Xidian

University were partially supported by the National Natural Science Foundation of China (No. 62302362) and the Key Research and Development Programs of Shaanxi under Grants S2024-YF-YBGY-1540. Yichang Xiong and Xiaokuan Zhang were partially supported by a seed grant from the CAHMP center at George Mason University.

REFERENCES

- [1] 3D-TOP-Event. 2023. Privacy Policies - No user data is collected from the app manufacturer. <https://3dtopevent.info/oculus/user-data.html>. Accessed on 2023/04/15.
- [2] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M Redmiles. 2018. Ethics emerging: the story of privacy and security perceptions in virtual reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. 427–442.
- [3] Adobe. 2023. Adobe Privacy Policy. <http://www.adobe.com/privacy.html>. Accessed on 2023/04/15.
- [4] Abdulrahman Alabduljabbar and David Mohaisen. 2022. Measuring the Privacy Dimension of Free Content Websites through Automated Privacy Policy Analysis and Annotation (*WWW '22*). Association for Computing Machinery, New York, NY, USA, 860–867. <https://doi.org/10.1145/3487553.3524663>
- [5] Benjamin Andow, Samin Yaseer Mahmud, Wenyu Wang, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Tao Xie. 2019. PoliLint: Investigating Internal Privacy Policy Contradictions on Google Play. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 585–602. <https://www.usenix.org/conference/usenixsecurity19/presentation/andow>
- [6] Benjamin Andow, Samin Yaseer Mahmud, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Serge Egelman. 2020. Actions Speak Louder than Words: Entity-Sensitive Privacy Policy and Data Flow Analysis with PoliCheck. In *29th USENIX Security Symposium (USENIX Security 20)*. 985–1002.
- [7] APPI. 2003. Act on the Protection of Personal Information. <https://www.cas.go.jp/jp/seisaku/hourei/data/APPI.pdf>. Accessed on 2023/04/15.
- [8] Noah Aporthe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 123–140. <https://www.usenix.org/conference/usenixsecurity19/presentation/aporthe>
- [9] Michael Backes, Sven Bugiel, Erik Derr, Patrick McDaniel, Damien Ocateau, and Sebastian Weisgerber. 2016. On Demystifying the Android Application Framework: Re-Visiting Android Permission Specification Analysis. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 1101–1118. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/backes_android
- [10] Jeremy Bailenson. 2018. Protecting nonverbal data tracked in virtual reality. *JAMA pediatrics* 172, 10 (2018), 905–906.
- [11] Behance. 2023. Behance. <https://www.behance.net/>. Accessed on 2023/04/15.
- [12] Benandow. 2023. HtmlToPlaintext. <https://github.com/benandow/HtmlToPlaintext>. Accessed on 2022/10/15.
- [13] Jasmine Bowers, Bradley Reaves, Imani N Sherman, Patrick Traynor, and Kevin Butler. 2017. Regulators, mount up! analysis of privacy policies for mobile money services. In *Thirteenth symposium on usable privacy and security (SOUPS 2017)*. USENIX Association, 97–114.
- [14] Ryan Browne. 2021. Bought your kid a VR headset for Christmas? You might end up regretting it. <https://www.cnbc.com/2021/12/31/bought-your-kid-a-vr-headset-for-christmas-you-might-regret-it.html>. Accessed on 2022/12/18.
- [15] Lauren E Buck and Bobby Bodenheimer. 2021. Privacy and personal space: Addressing interactions and interaction data as a privacy concern. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 399–400.
- [16] Duc Bui, Kang G Shin, Jong-Min Choi, and Junbum Shin. 2021. Automated Extraction and Presentation of Data Practices in Privacy Policies. *Proc. Priv. Enhancing Technol.* 2021, 2 (2021), 88–110.
- [17] Duc Bui, Brian Tang, and Kang G Shin. 2023. Detection of Inconsistencies in Privacy Practices of Browser Extensions. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2780–2798.
- [18] Kent Bye, Diane Hosfelt, Sam Chase, Matt Miesnieks, and Taylor Beck. 2019. The ethical and privacy implications of mixed reality. In *ACM SIGGRAPH 2019 Panels*. 1–2.
- [19] Mingliang Cao, Tianhua Xie, and Zebin Chen. 2019. Wearable sensors and equipment in VR games: a review. *Transactions on Edutainment XV* (2019), 3–12.
- [20] CCPA. 2020. California Consumer Privacy Act. <https://oag.ca.gov/privacy/ccpa>. Accessed on 2023/04/15.
- [21] VR Chat. 2023. PRIVACY POLICY. <https://hello.vrchat.com/privacy>. Accessed on 2023/04/15.
- [22] CPRA. 2023. The California Privacy Rights Act of 2020. <https://theccpra.org/>. Accessed on 2023/04/15.

- [23] Hao Cui, Rahmadi Trimmananda, Athina Markopoulou, and Scott Jordan. 2023. PoliGraph: Automated Privacy Policy Analysis using Knowledge Graphs. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 1037–1054. <https://www.usenix.org/conference/usenixsecurity23/presentation/cui>
- [24] GAYATRI MURTHY DAVID MEDINE. 2019. Nobody Reads Privacy Policies: Why We Need to Go Beyond Consent to Ensure Data Privacy. <https://nextbillion.net/beyond-consent-for-data-privacy/>. Accessed on 2023/04/15.
- [25] Jaybie A De Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2019. Security and privacy approaches in mixed reality: A literature survey. *ACM Computing Surveys (CSUR)* 52, 6 (2019), 1–37.
- [26] FTC. 2023. Children’s Online Privacy Protection Rule: A Six-Step Compliance Plan for Your Business. <https://www.ftc.gov/business-guidance/resources/childrens-online-privacy-protection-rule-six-step-compliance-plan-your-business>. Accessed on 2023/09/15.
- [27] GDPR. 2018. General Data Protection Regulation. <https://gdpr-info.eu/>. Accessed on 2023/04/15.
- [28] Thomas Germain. 2022. Meta’s New Headset Will Track Your Eyes for Targeted Ads. <https://gizmodo.com/meta-quest-pro-vr-headset-track-eyes-ads-facebook-1849654424>. Accessed on 2023/04/15.
- [29] Google. 2023. Introducing Checks. <https://checks.google.com/>. Accessed on 2023/04/15.
- [30] Sonu Gupta, Ellen Poplavska, Nora O’Toole, Siddhant Arora, Thomas Norton, Norman Sadeh, and Shomir Wilson. 2022. Creation and Analysis of an International Corpus of Privacy Laws. arXiv:2206.14169 [cs.CL]
- [31] Jassim Happa, Anthony Steed, and Mashhuda Glencross. 2021. Privacy-certification standards for extended-reality devices and services. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 397–398.
- [32] Hamza Harkous, Kassem Fawaz, Rémi Leuret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *Proceedings of the 27th USENIX Conference on Security Symposium (Baltimore, MD, USA) (SEC’18)*. USENIX Association, USA, 531–548.
- [33] DAVID HEANEY. 2022. IDC Estimates Quest 2 Has Sold Almost 15 Million Units. <https://www.uploadvr.com/quest-2-sold-almost-15-million-idc/>. Accessed on 2023/04/15.
- [34] Jane Henriksen-Bulmer and Sheridan Jeary. 2016. Re-identification attacks—A systematic literature review. *International Journal of Information Management* 36, 6 (2016), 1184–1192.
- [35] Apple Inc. 2024. Apple Vision Pro. <https://www.apple.com/apple-vision-pro/>.
- [36] Intel. 2023. Intel Privacy Notice. <https://www.intel.com/content/www/us/en/privacy/intel-privacy-notice.html>. Accessed on 2022/12/18.
- [37] Outlaw Jessica, Persky Susan, Bailensen Jeremy, Bye Kent, and Rosedale Philip. 2018. Industry review boards are needed to protect VR user privacy. <https://www.extendedmind.io/xr-privacy-summit>.
- [38] Jingdong Jia and Wenchao Chen. 2017. The ethical dilemmas of virtual reality application in entertainment. In *2017 IEEE International Conference on Computational Science and Engineering (CSE) and IEEE International Conference on Embedded and Ubiquitous Computing (EUC)*, Vol. 1. IEEE, 696–699.
- [39] LGPD. 2022. Brazilian General Data Protection Law. <https://lgpd-brazil.info/>. Accessed on 2023/04/15.
- [40] Youqun Li, Yichi Zhang, Haojin Zhu, and Suguo Du. 2021. Toward Automatically Generating Privacy Policy for Smart Home Apps. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 1–7.
- [41] Sunil Manandhar, Kaushal Kafle, Benjamin Andow, Kapil Singh, and Adwait Nadkarni. 2022. Smart Home Privacy Policies Demystified: A Study of Availability, Content, and Coverage. In *31st USENIX Security Symposium (USENIX Security 22)*. 3521–3538.
- [42] Rachel McAmis and Tadayoshi Kohno. 2023. The Writing on the Wall and 3D Digital Twins: Personal Information in (not so) Private Real Estate. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 2169–2186. <https://www.usenix.org/conference/usenixsecurity23/presentation/mcamis>
- [43] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.
- [44] Meta. 2022. Meta Connect 2022: Meta Quest Pro, More Social VR and a Look Into the Future. <https://about.fb.com/news/2022/10/meta-quest-pro-social-vr-connect-2022/>. Accessed on 2023/04/15.
- [45] Meta. 2023. Complete a Data Use Checkup. <https://developer.oculus.com/resources/publish-data-use/>. Accessed on 2023/04/15.
- [46] Meta. 2023. Get a Front Row Seat to NBA Games on Meta Quest. <https://about.fb.com/news/2023/01/nba-games-in-vr-on-quest/>. Accessed on 2023/09/15.
- [47] Meta. 2023. Meta Quest Virtual Reality Check (VRC) Guidelines. <https://developer.oculus.com/resources/publish-quest-req/>. Accessed on 2023/04/15.
- [48] Meta. 2023. Oculus Privacy Policy. <https://www.oculus.com/legal/privacy-policy/>. Accessed on 2023/04/15.
- [49] Meta. 2023. Privacy Policy Requirements. <https://developer.oculus.com/policy/privacy-policy/>. Accessed on 2023/04/15.
- [50] Cade Metz. 2023. A New Way for Therapists to Get Inside Heads: Virtual Reality. <https://www.nytimes.com/2017/07/30/technology/virtual-reality-limbic-mental-health.html>. Accessed on 2023/09/15.
- [51] Microsoft. 2023. Create an app submission for your app. <https://learn.microsoft.com/en-us/windows/apps/publish/publish-your-app/create-app-submission?pivot=store-installer-msix>. Accessed on 2023/04/15.
- [52] Microsoft. 2023. Microsoft privacy policy. <https://privacy.microsoft.com/en-us/privacystatement>. Accessed on 2023/04/15.
- [53] Mikeage. 2023. GetSidequestURL. https://github.com/mikeage/get_sidequest_urls. Accessed on 2023/09/21.
- [54] Vivek Nair, Gonzalo Munilla Garrido, and Dawn Song. 2022. Exploring the Unprecedented Privacy Risks of the Metaverse. arXiv:2207.13176 [cs.CR]
- [55] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O’Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 895–910. <https://www.usenix.org/conference/usenixsecurity23/presentation/nair-identification>
- [56] DEF CON CHINA PARTY. 2023. Enter the first virtual reality DEF CON Experience. https://defcon.org/html/defcon-china-party/dc-china-party.html?fbclid=IwAR1Rda7vJIE430ib-stB SRCxFx DN-te1p_Vm1RBbgGPu3apnguhBZpfVnVo. Accessed on 2023/09/15.
- [57] Alfredo J Perez, Sherali Zeadally, and Jonathan Cochran. 2018. A review and an empirical analysis of privacy policy and notices for consumer Internet of things. *Security and Privacy* 1, 3 (2018), e15.
- [58] PIPEDA. 2022. The Personal Information Protection and Electronic Documents Act. <https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/>. Accessed on 2023/04/15.
- [59] PIPL. 2021. Personal Information Protection Law of the People’s Republic of China. http://en.npc.gov.cn.cdurl.cn/2021-12/29/c_694559.htm. Accessed on 2023/04/15.
- [60] Simon Read. 2023. How many consumers are shopping in virtual reality and what can it offer them? <https://www.weforum.org/agenda/2022/08/virtual-reality-shopping-retail/>. Accessed on 2023/09/15.
- [61] Nelson Reed. 2022. What Are Kids Doing in the Metaverse? <https://www.common sense media.org/kids-action/articles/what-are-kids-doing-in-the-metaverse>. Accessed on 2023/04/15.
- [62] RoadToVR. 2023. Varjo Signs “multi-million dollar” Deal to Provide Headsets for Army Training Systems. <https://www.roadtovr.com/varjo-rvct-ste-deal-cole-engineering-flight-sim-training/>.
- [63] Web Scraper. 2022. Web Scraper. <https://www.webscraper.io/>. Accessed on 2022/12/16.
- [64] Zimmeck Sebastian, Wang Ziqi, Zou Lieyong, Iyengar Roger, Liu Bin, Schaub Florian, Wilson Shomir, Sadeh Norman, Steven M. Bellovin, and Reidenberg Joel. 2017. *Automated Analysis of Privacy Requirements for Mobile Apps*. Korea Society of Internet Information. <https://doi.org/10.14722/nds.2017.23034>
- [65] Selenium. 2022. Selenium. <https://pypi.org/project/selenium/>. Accessed on 2022/12/16.
- [66] Laura Shipp and Jorge Blasco. 2020. How private is your period?: A systematic analysis of menstrual app privacy policies. *Proc. Priv. Enhancing Technol.* 2020, 4 (2020), 491–510.
- [67] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (appropriate) flow: A contextual integrity approach to privacy policy analysis. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 162–170.
- [68] SIDEQUEST. 2023. Disclaimer for Third-Party Content. <https://sidequestvr.com/terms>. Accessed on 2023/09/15.
- [69] Mukund Srinath, Shomir Wilson, and C Lee Giles. 2021. Privacy at Scale: Introducing the PrivaSeer Corpus of Web Privacy Policies. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 6829–6839. <https://doi.org/10.18653/v1/2021.acl-long.532>
- [70] Welderufael B. Tesfay, Peter Hofmann, Toru Nakamura, Shinsaku Kiyomoto, and Jetzabel Serna. 2018. I Read but Don’t Agree: Privacy Policy Benchmarking Using Machine Learning and the EU GDPR. In *Companion Proceedings of the The Web Conference 2018 (Lyon, France)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 163–166. <https://doi.org/10.1145/3184558.3186969>
- [71] Rahmadi Trimmananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing Network Traffic and Privacy Policies in Oculus VR. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 3789–3806. <https://www.usenix.org/conference/usenixsecurity22/presentation/trimananda>
- [72] Columbia University. 2023. Virtual Reality Training for Doctors Made Real by Columbia’s Clinical Innovation Lab. <https://www.cuimc.columbia.edu/news/>

- virtual-reality-training-doctors-made-real-columbias-clinical-innovation-lab. Accessed on 2023/09/18.
- [73] urllib. 2023. urllib. <https://docs.python.org/3/library/urllib.html>. Accessed on 2023/04/15.
- [74] LA BIENNALE DI VENEZIA. 2023. VENICE IMMERSIVE. <https://www.labiennale.org/en/cinema/2023/venice-immersive-0>. Accessed on 2023/09/15.
- [75] Viveport. 2023. Distribution & Price. <https://developer.vive.com/resources/viveport/store-guide/store-submission-guide/english/submitting-your-viveopenvr-content/distribution-price/>. Accessed on 2023/04/15.
- [76] IMMERSION VR. 2023. What is virtual reality in travel? <https://immersionvr.co.uk/about-360vr/vr-for-tourism/>. Accessed on 2023/09/15.
- [77] VRChat. 2023. Owner estimations. <https://steamdb.info/app/438100/charts/>. Accessed on 2023/04/15.
- [78] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N. Cameron Russell, Thomas B. Norton, Eduard Hovy, Joel Reidenberg, and Norman Sadeh. 2016. The Creation and Analysis of a Website Privacy Policy Corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1330–1340. <https://doi.org/10.18653/v1/P16-1126>
- [79] Weltsehn XR. 2022. XR Hardware Disassembly and BOM Cost Report for Meta Quest Pro. <https://vr.xr-expert.com/meta-quest-pro-review/>. Accessed on 2023/04/15.
- [80] Le Yu, Xiapu Luo, Jiachi Chen, Hao Zhou, Tao Zhang, Henry Chang, and Hareton K. N. Leung. 2021. PPChecker: Towards Accessing the Trustworthiness of Android Apps' Privacy Policies. *IEEE Transactions on Software Engineering* 47, 2 (2021), 221–242. <https://doi.org/10.1109/TSE.2018.2886875>
- [81] Le Yu, Tao Zhang, Xiapu Luo, and Lei Xue. 2015. AutoPPG: Towards Automatic Generation of Privacy Policy for Android Applications. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices* (Denver, Colorado, USA) (SPSM '15). Association for Computing Machinery, New York, NY, USA, 39–50. <https://doi.org/10.1145/2808117.2808125>
- [82] Le Yu, Tao Zhang, Xiapu Luo, Lei Xue, and Henry Chang. 2016. Toward automatically generating privacy policy for android apps. *IEEE Transactions on Information Forensics and Security* 12, 4 (2016), 865–880.
- [83] Xiao Yu, Yiyu Yang, Wenjie Wang, and Yuqing Zhang. 2021. Whether the sensitive information statement of the IoT privacy policy is consistent with the actual behavior. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 85–92.
- [84] Lu Zhou, Chengyongxiao Wei, Tong Zhu, Guoxing Chen, Xiaokuan Zhang, Suguo Du, Hui Cao, and Haojin Zhu. 2023. POLICYCOMP: Counterpart Comparison of Privacy Policies Uncovers Overbroad Personal Data Collection Practices. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 1073–1090.
- [85] Sebastian Zimmeck, Rafael Goldstein, and David Baraka. 2021. PrivacyFlash Pro: Automating Privacy Policy Generation for Mobile Apps. In *Network and Distributed Systems Security (NDSS) Symposium*.

A VR SYNONYMS AND ONTOLOGIES.

To alleviate the manual burden of checking a large number of un-terminologized phrases, we propose an approach based on the insight that synonyms have similar semantics and are therefore clustered in the embedding space.

To process data object phrases, we first utilize a BERT-based sentence embedding model to map all phrases to the semantic embedding space. During this phase, any phrases that are within a threshold distance (0.8 in our study, i.e., the median context similarity of the OVRSeen synonyms file) are added to synonym lists. For remaining unterminologized phrases, we iteratively spot new clusters in embedding space and determine whether they can be included in the VR data ontology and where they should be placed. During this process, we assume a broadly inclusive position for what personal information counts as PII only if they do not claim to be non-PII, as prior work [34] has shown identity can be reconstructed from a variety of information types. As a result, we obtain an extended VR data ontology with 107 nodes and corresponding synonym lists containing 8,042 distinct data object phrases (See Figure 11(a)).

Algorithm 1: LowerBound

Input: All data objects D in extracted CUS tuples from a privacy policy and data ontology O (whose node set is V)

Output: Lower bounds LB of D .

// Calculate NodeGranularity (NG) of each node and store them accordingly.

$G \leftarrow []$;

$maxNG \leftarrow 0$;

for $node\ n\ in\ V$ **do**

$ng = NG(O, n)$;

 update $maxNG$;

 append n to G_{ng} ;

end

// Initialize Lower bounds as all leaves in D

$LB \leftarrow G_0$;

// Iterate each node by their NG score

for $i = 1$ to $maxNG$ **do**

$NodesOfInterest = D \cap G_i$;

$NodesToDelete = \{\}$;

for $node\ u\ in\ NodesOfInterest$ **do**

for $node\ v\ in\ LB$ **do**

if $HasPath(G, u, v)$ **then**

 add $node\ u$ to $NodesToDelete$;

break;

end

end

end

$LB = (LB \cup NodesOfInterest) \setminus NodesToDelete$;

end

Algorithm 2: UpperBound

Input: All data objects D in extracted CUS tuples from a privacy policy, data ontology O

Output: Upper bounds UB of D .

$UpperBounds \leftarrow D$;

for $node\ n\ in\ D$ **do**

$UB = UB \cup Descendants(O, n)$;

end

When dealing with entity phrases, we consider two situations: entity category and company name. Entity categories such as *payment processor* describe the functionality of an entity and have semantic meaning. For instance, *credit card company* and *billing company* are close in embedding space and can be considered synonyms for *payment processor*. On the other hand, company names like *checkout* and *braintree*, though both are payment processors, share no semantic similarity. Therefore, we identify synonyms of the company names through keyword matching since their names are context-free and remain identical. When handling implicit first-party entities (i.e., using the company name instead of *we*), we follow the guidelines outlined in [71], and also take into account (1) the app name, (2) publisher or developer name, and (3) domain names found in both homepage links and privacy policy links. In total, we obtain an extended VR entity ontology containing 117 nodes along with synonym lists containing 1,663 distinct entity phrases for each category of entities (See Figure 11(b)).

B LEGAL BASIS OF FIVE CRITERIA

Some legal policies (e.g., GDPR, CCPA, PIPL) related to our vetting criteria are listed in Table 10.

Table 10: Legal articles related to vetting 5 criteria across various regions.

Requirements	GDPR	CCPA	PIPL	
PP	C1 Availability	5.1(a,b), 12.1	1798.100	1
	User Choice	5.1(a), 6, 7, 8, 9, 12-22, 49	1798.120, 999.306, 999.308(c)(3)	15, 24, 44
	Data Collection	5.1(a-c), 4-11, 19, 24-25, 28-30, 33-29, 44-49	1798.100, 1798.115, 999.305, 999.308	17, 23
	User Rights	5.1(a,d), 12-22	1798.105, 1798.106, 1798.110, 999.308(c)(2)	45-47
	C2 Data Retention	5.1(e), 5, 25, 30	1798.125	17
	Data Security	5.1(f), 6, 32-34	999.313, 999.326	57
	Policy Change	5.1(a)	-	17
CUS	Spec. Audience	5.1(a), 8	999.330-332	28,31
	C3 Granularity	5.1(a-b), 24-29	1798.100	7
	C4 Minimization	5.1(c)	1798.140(e)	6
C5 Consistency	5.1(a-e)	1798.100	7	

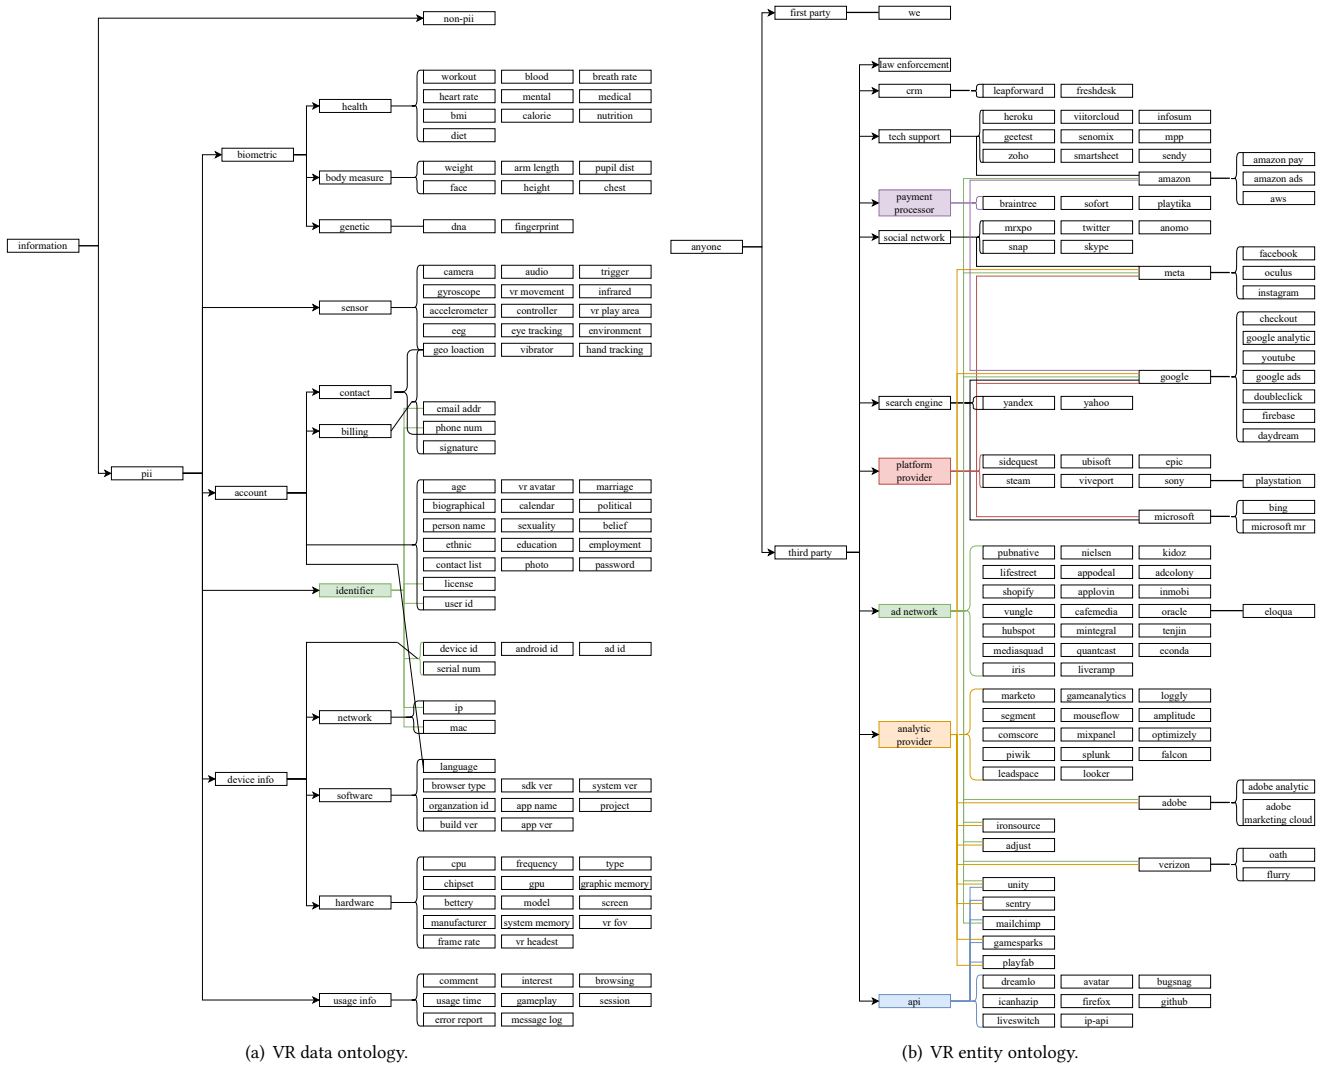


Figure 11: Extended VR data ontology and entity ontology. To be concise, some arrows are omitted and the left bracket indicates edges from the parent node to all nodes inside.

C EXAMPLES OF VR APPS VIOLATING VETTING CRITERIA

Tables 11, 12, and 13 show some typical VR apps which violate the criteria of availability, structural completeness, and minimization requirement.

