


Please cite the Published Version

Bolton, Tom, Dargahi, Tooska , Belguith, Sana and Maple, Carsten (2023) PrivExtractor: toward redressing the imbalance of understanding between virtual assistant users and vendors. ACM Transactions on Privacy and Security, 26 (3). 31 ISSN 2471-2566

DOI: <https://doi.org/10.1145/3588770>

Publisher: Association for Computing Machinery (ACM)

Version: Accepted Version

Downloaded from: <https://e-space.mmu.ac.uk/633048/>

Usage rights:  In Copyright

Additional Information: © Authors 2023. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Transactions on Privacy and Security, <http://dx.doi.org/10.1145/3588770>.

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

Privextractor: Towards Redressing the Imbalance of Understanding Between Virtual Assistant Users and Vendors

Thomas Bolton

School of Science, Engineering and Environment, University of Salford, UK, t.j.e.bolton@edu.salford.ac.uk

Tooska Dargahi

Department of Computing and Mathematics, Manchester Metropolitan University, UK, t.dargahi@mmu.ac.uk

Sana Belguith

Department of Computer Science, University of Bristol, UK, sana.belguith@bristol.ac.uk

Carsten Maple

Warwick Manufacturing Group (WMG), University of Warwick, UK, Maple, CM@warwick.ac.uk

The use of voice-controlled virtual assistants (VAs) is significant, and user numbers increase every year. Extensive use of VAs has provided the large, cash-rich technology companies who sell them with another way of consuming users' data, providing a lucrative revenue stream. Whilst these companies are legally obliged to treat users' information 'fairly and responsibly', artificial intelligence techniques used to process data have become incredibly sophisticated, leading to users' concerns that a lack of clarity is making it hard to understand the nature and scope of data collection and use.

There has been little work undertaken on a self-contained user awareness tool targeting VAs. Privextractor, a novel web-based awareness dashboard for VA users, intends to redress this imbalance of understanding between the data 'processors' and the user. It aims to achieve this using the four largest VA vendors as a case study and providing a comparison function that examines the four companies' privacy practices and their compliance with data protection law.

As a result of this research, we conclude that the companies studied are largely compliant with the law, as expected. However, the user remains disadvantaged due to the ineffectiveness of current data regulation that does not oblige the companies to fully and transparently disclose how and when they use, share, or profit from the data. Furthermore, the software tool developed during the research is, we believe, the first that is capable of a comparative analysis of VA privacy with a visual demonstration to increase ease of understanding for the user.

CCS CONCEPTS • Security and privacy • Human and societal aspects of security and privacy • Usability in security and privacy

1 INTRODUCTION

Cash-rich, monolithic private-sector technology companies are significant consumers of personal information with their primary goal being revenues achieved from targeted advertising. Four of the world's five richest corporations are Microsoft, Apple, Amazon, and Google; as of 2021, the four have a combined market capitalisation of US\$7.47 trillion [1]. All four hold vast quantities of data relating to individuals that they use to sell targeted advertising [2] [3] [4] [5]. Google's total company-wide revenue in 2021 was US\$257 billion, of which the majority, US\$209 billion, came via its Google Ads platform [6]. These four companies are not alone in marketing a voice assistant (VA), or in using personal data for advertising purposes – far from it. They are, however, the biggest technology companies in the world in monetary terms. For this reason, and to make the research more manageable, Microsoft, Apple, Amazon, and Google will be the subjects of the case study in this paper.

George Orwell's future-dystopian fictional novel *Nineteen Eighty-Four*, published in 1949, refers to an electronic device called the 'Telescreen' that has the ability to see, hear, and broadcast; in Orwell's fictional world, the Telescreen is forced into every citizen's home by law. Today, the VA, a centrally-controlled hearing and broadcasting device, is marketed as a must-have lifestyle accessory in a range of prices and designs [7] for which people voluntarily pay. The VA forms part of the private sector's move towards large-scale collection and processing of personal data.

The complex and lucrative business of brokering online advertising relies on data that describes the user and their preferences. One common source of this data is a user's web browsing history – their 'click behaviour' [8]. This data has traditionally been collected via a user's keyboard input; however, companies have recently begun to use newly emerging computing devices, equipped with microphones to enable data capture in the form of a user's speech, to harness additional forms of information.

VAs such as Apple's Siri and Amazon's Alexa are software applications with which users can interact verbally, almost conversationally. In return, the assistant can provide information, or can interact with devices around the home to which it is connected – for example, to play music or switch off a light. It is important to understand, however, that regardless of the device upon which the VA is installed, that device is simply an endpoint – the majority of the work in servicing the user's voice command is carried out on the provider's servers [9]. Section 3.1: Forensic Recovery outlines some findings that show artefacts recovered from both Amazon's Alexa and Microsoft's Cortana – both in text form, and in the form of recorded audio. These artefacts have been transmitted to, and stored on, the vendors' cloud backend. This transmission and storage constitute what GDPR defines as 'processing' and, as such, is subject to that data law.

It is clear that in an industry that is home to large cash-rich private companies, whose profit is largely (or even partly) reliant on the gathering of data from individuals, the balance of power lies with those companies and not with the end user of the products. Understanding the mechanics which underpin the process of bidding for advertising requires a good deal of knowledge of computer science; even understanding the impact of the mechanics is not straightforward. Smit et al. conducted a study in which participants were questioned on their understanding of online advertising; 41.1% of participants in the survey believed that *"When I visit a website, I see the same ads as someone else visiting that website"*, contrasting with the 82.5% of participants who believed that *"Your browsing history determines which ads you are going to see during your next visit."* [10]. This disconnect suggests that, whilst users are aware that their browsing history is being mined, they do not necessarily understand the impact this use of their data has on what they see when viewing online advertising.

VA users are becoming concerned about the lack of clarity that makes it hard to understand the nature and scope of data collection [11]. Any user who wishes to better comprehend how their data is collected and used is reliant on the vendors to explain this. However, whilst there is data law such as the General Data Protection Regulation (GDPR) governing the behaviour and responsibilities of the vendors, it has been shown by Linden et al. that “*many [vendors’ privacy] policies still do not meet several key GDPR requirements or their improved coverage comes with reduced specificity*” [12].

1.1 Related Work

A VA is a software application; more accurately, it is a whole series of connected software applications. A VA’s client can take many forms. Software translations (or ‘ports’) of commercial VA clients such as Alexa are available for smartphones, tablets, televisions, TV ‘sticks’ such as Google’s Chromecast, video games consoles, and dedicated smart speakers – to name a selection. 38.2% of adults in the United Kingdom have adopted a smart speaker in the home – a growth of 24.5% during the global COVID-19 pandemic [13]. In 2019, an estimated 3.25 billion VAs were being used globally. Forecasts further estimate that by 2023 this number might hit eight billion globally at which point VAs will outnumber humans [14].

Having thoroughly researched the problem, we could find no equivalent privacy dashboards which examine and compare VAs in terms of privacy and data law compliance in the way that Privextractor does. Tools which address privacy on mobile devices exist, as do tools which help a user manage the settings of their VA devices, but none is comparable directly to PrivExtractor.

There does appear to be a requirement for such a tool: Sharma et al. made a study of VA users and their perceptions towards Google’s Assistant dashboard and found some concerns [15]. 38.7% of users were unaware that Google collects audio recordings; when shown transcripts of these interactions with Google, the authors found that the participants would be ‘uncomfortable’ sharing around 18% of their individual conversations with the company. In studying the vendors, Liao et al. [11] applied a comprehensive, quantitative technical approach to analysing the privacy policies of two VAs – Google Assistant and Amazon Alexa – along with the policies ascribed to the software add-ins for each platform. Using data mining and machine learning techniques, the authors analysed the vendors’ data practices and found some alarming results: not only were there many incorrect privacy policy URLs and broken links, but some of Amazon and Google’s voice apps violated their own policy.

Zibuschka et al. have identified these privacy concerns and presented ENTOURAGE – a ‘privacy and security reference architecture’; part of the authors’ system presents a dashboard to the user through which they may control the privacy settings of their VA and manage data [16]. ENTOURAGE does not, however, offer an insight into data law and privacy compliance of similar devices competing in the VA marketplace.

Privacy Flag, a European research project co-funded by the European Commission and the Swiss State Secretariat for Education, Research and Innovation, has produced a smartphone application that aims to inform users of privacy-related risks emanating from applications installed on the same device [17]. Privacy Flag’s backend system gathers inputs from ‘technical enablers’ and via crowdsourcing in order to inform a user that an application installed on their device is, or is not, privacy friendly. This is an interesting project with much scope for restoring the imbalance between users and the providers of the applications. Currently, the application exists for Android devices only – porting the work to other mobile

operating systems would enable a much larger user base to benefit. The work undertaken here mirrors, to some degree, what we would like to achieve with VAs and VA devices - whilst it would be more difficult to create an app native to the devices, which lack touch and screen interfaces, there are many parallels with our goals.

Finding that privacy policies are “...*excessively long and difficult to follow*”, Harkous et al. created Polisis – a framework that uses machine learning-based analysis of privacy policies to divide a policy into smaller, self-contained fragments for easier digestion by the user [18]. The authors found that their application was able to apply icons to the privacy policy such as ‘Precise location’ and ‘Data retention’ to an accuracy of 88.4%. As a fundamental part of PrivExtractor is the analysis of VAs’ privacy policies, this research is of interest as a means to avoid the manual analysis of the relevant policies.

1.1.1 Data and the Law

To regulate the privacy of its citizens whose details were increasingly being recorded in government databases, Sweden introduced the Data Act in 1973; this was not the first data law, but was the world’s first of its kind to apply to an entire nation [19]. The United Kingdom took until 1987 to introduce its law – the Data Protection Act (DPA) – which would be enforced by the newly-assembled Information Commissioner’s Office (ICO) [20].

The General Data Protection Regulation (GDPR) comprises data security and privacy law and came into effect on 25 May 2018; it is claimed to be the “*toughest privacy and security law in the world*” [21]. It was drafted and passed by the European Union (EU) and imposes legal responsibilities on organisations anywhere in the world so long as they target or collect the data of people resident in the EU. The regulation, in its current state, comprises 99 articles [22]. It was the introduction of GDPR that led to a new DPA being made law in 2018 [20]. It should be noted that GDPR applies by itself and does not require national implementation. However, the DPA is the benchmark for UK data protection law and required changing to better reflect the comprehensive new regulation laid down in GDPR.

Following the UK’s formal exit from the EU on 31 December 2020, most of the EU GDPR was retained in UK law; this retained GDPR is known as the “UK GDPR” [23]. The DPA is still the UK’s primary data protection law; the UK GDPR sits alongside the DPA and applies to controllers and processors based outside the UK if they should offer goods or services to individuals in the UK, or monitor the behaviour of UK individuals.

In GDPR, an adult is any data subject aged 16 years or over; section 9 of the DPA lowers this to the minimum allowed under GDPR – 13 years [24]. To bridge that gap, the ICO has written a code of conduct for organisations dealing with the data of children and young adults. The ICO’s age-appropriate design is a code of practice for online services that came into force on 2 September 2020. The code explains how organisations can ensure their online services appropriately safeguard children’s personal data. It is intended that organisations can use it to demonstrate that they comply with GDPR [25]. It is important to note that, whilst the GDPR and DPA are generally applicable laws, the ICO’s Age-appropriate design code is limited in its applicability and does not have the same legislative significance.

1.1.2 User Perceptions

Lau et al. [26] found that people who choose to adopt a VA have worries that differ from those who do not. Those who refuse to see the purpose of such a device are more likely to hold privacy concerns. It is these users who are, for example, “*deeply uncomfortable with the idea of a ‘microphone-based’ device that a speaker company or an ‘other’ with malicious*

intent could ostensibly use to listen in on their homes”. Users who are keen to use VAs hold fewer concerns; this lack of worry is rationalised with the belief that the vendor can be trusted, and that it would be impossible for an unauthorised individual to access their data.

In a control group of users, acceptance factors of various VAs were considered by Burbach et al [27]. Using a choice-based conjoint analysis with three attributes - natural language processing (NLP) performance, price, and privacy - the authors found that privacy is the chief concern among users. These findings appear to loosely tally with those of Lau et al [26]. However, the surveys used were quite different in design, and the primary goal of the two studies was also distinct. Burbach et al.’s study was comprehensive in its design; the authors divided the survey participants into groups according to their chief concern. Only one group, named the ‘Thrifs’ by the authors, were concerned by the price of the VA as opposed to its privacy, or potential lack thereof. This segment, however, formed only 18% of the total user group who were, overwhelmingly, concerned about privacy more than cost or NLP performance. Combining the attributes of cost and privacy, Ebbers et al. analysed user preferences of key attributes of VA privacy features – the amount of personal data shown to users, the explainability of the VA’s decisions, and the gamification level of the UI [28]. A key finding showed that 56.4% of participants would be willing to pay for privacy features; these users were young and concerned about privacy.

Again taking a technical approach, this time combined with user education, Seymour et al. [29] developed the software tool, Aretha. Designed to demonstrate to a user both the data coming in and out of the home and the ramifications of this, Aretha was deployed in three users’ homes and the users’ behaviour was observed. One finding, in particular, is interesting for our study: *“The lack of engagement with the firewall was instructive in its own way; while most participants found it difficult to use effectively, due to having already observed, interpreted, and understood the underlying behaviour of their devices they appeared better able to adapt, invent, or imagine other protective mechanisms, tools, and strategies.”* This observation suggests that, when armed with clear information as to how their data is treated, users feel more empowered to control what is shared with others.

1.1.3 Security

We can see then that privacy, and the security necessary to ensure that VA interactions and information remain private, are important to users. The field of usable security examines, amongst other aspects, how security is traded with usability; a review by Lennartsson et al. finds that *“Usability is hampered when users’ primary tasks are disturbed.”* [30]. Further findings, that *“Necessary security actions should be arranged in ways that minimize interruptions”* and *“compelling users to remember passwords repeatedly interrupts other tasks as enforced context switches may cause confusion”* are of interest when viewing security implications of a VA – a tool that is, by design, friction-free in its use of a voice interface.

Yan et al. conducted a comprehensive survey on VA cyber attacks and countermeasures that users might employ [31]. Perhaps counterintuitive to Lennartsson et al.’s findings on usable security, Yan et al. recommend to users that they might avoid using their VA for bank account management or home unlocking, avoid leaving the VA unattended, disallow wake-word detection, or disable the VA when the device is locked. These are all solid suggestions; however, they tend to fly in the face of the VA’s unique, straightforward mode of operation. Interestingly, one finding from Yan et al. is that *“Despite that it is the users who actually suffer from the security consequences, there are relatively few [things] they can do to prevent the attacks.”*

1.1.4 Children and VAs

McReynolds et al. researched the comparatively novel area of connected toys, such as Jibo and Hello Barbie [32]. These were studied in parallel with ‘adult’ VAs to answer a number of questions – chiefly whether a child relates to a VA in the same way it would the toys. The authors found that five of the nine parents who took part in the study – when asked if their children interacted with ‘adult’ VAs - “...explicitly observed that *Dino* [toy] was similar to *Siri* and other artificial intelligence voice recognition systems”. The relationship children have with VAs was the subject of a work by Girouard-Hallam et al. who made a study revealing children’s perceptions of a VA and, in particular, whether the child thought of the VA as having mental, social, and moral attributes [33]. This was a comprehensive study with some revealing findings around how children perceive and interact with VAs; 65% of participants responded ‘yes’ to the question ‘Can [device] be your friend?’. In summary, the authors note that “*Children’s beliefs about their social, or perhaps parasocial, relationships with voice assistants may also influence their understanding of cybersafety. Believing that a voice assistant could keep a secret, and that it is, at least in part, amoral entity, may contribute to young children oversharing information with internet-based devices*”; the differences in the way in which VA vendors treat children and their data, and an adult’s, will be part of this study.

1.1.5 Technology and Forensics

Javed et al. [34] conducted an in-depth study of what Alexa is listening to. Disputing Amazon’s claim that until the wake word is used no recording will take place, the study found that the VA was indeed recording: 91% of a control group of users had experienced such an unwanted episode. After investigation, it was discovered that passive sounds – radio, television, background noise – were recorded in the majority of these cases. More seriously, and representing a privacy concern, were the recordings made of sensitive information experienced by 29.2% of the study group.

There exists little documentation of the finer details of how VAs communicate with their cloud services; Ford et al. undertook a study of Amazon’s Alexa and its voice streaming network traffic, ostensibly to discover if VA devices were recording and streaming conversation without being prompted by the user [35]. Finding that Alexa’s internet traffic uses Transport Layer Security (TLS) for its communications with the cloud service, and not having a key with which to decrypt the traffic, the authors were forced to resort to observing patterns in the quantity of data that is exchanged between the device and its cloud platform in order to make any useful analysis.

Akinbi et al. attempted to recover forensic data from an Android smartphone running both Google Home and Google Assistant that was also used to control a Google Nest device [36]. The authors found useful forensic artefacts on the device, along with a chronology of voice interactions. For the purposes of this study, the most interesting finding was the ability to recover copies of past conversations from the user’s Google cloud service account. One of the more comprehensive studies made of VA forensics was that by Chung et al [37]. The authors made a thorough examination of Amazon’s Alexa ecosystem and were able to extract artefacts from both device and cloud, and develop a web-based dashboard to display the information in a user-friendly manner. The cloud artefact extraction is of particular interest as the exercise revealed an undocumented Web API that could be queried in order to retrieve data pertaining to both the user’s account and their interactions with Alexa.

Microsoft’s VA, Cortana, was the subject of a study made in 2017 by Singh et al. [38]. The authors were able to extract and examine forensic artefacts from local database storage and wrote Python scripts to simplify this process for future

investigators. Possibly due to the ‘walled garden’ nature of Apple’s mobile operating systems, there are fewer studies available that focus on Siri. One such was made by Horsman in 2019 [39] in which the author noted the information that could be extracted from Siri on a locked iOS device using carefully crafted voice interactions.

1.2 Research Questions

There has been little work undertaken in the development of a self-contained user awareness tool targeting virtual assistants. This problematic imbalance of understanding between the vendors and the end user, and the lack of access to clear information regarding the vendors’ adherence to data law, is what this research seeks to address. To this end, the following research questions (RQ) are asked:

RQ1: If a user of a voice assistant wishes to know the extent to which their personal information is harvested by the vendor of their chosen VA, does that vendor clearly and unequivocally state the exact nature of the information that they collect, how securely they keep that data, what they are doing with it, and for how long they keep it?

RQ2: Using the UK and European data law as a basis – GDPR, the Data Protection Act 2018 (DPA) and the Information Commission Office’s (ICO) age-appropriate code of conduct – can it be demonstrated to a user where the data collection practices of the VA vendor conflict with the law?

1.3 Contributions

This paper makes the following contributions:

- We systematically analyse the privacy policies of four major VA vendors, as a case study, to determine if those policies comply fairly with the UK and European data laws. We find that the problems are twofold: the vendors, whilst ostensibly complying with data laws such as GDPR, give little information to enable the user to see exactly how their data is manipulated. However, the blame for this cannot be placed solely upon the vendors: our analysis demonstrates that the data law itself offers insufficient requirements for specific transparency on behalf of the vendor.
- We use this analysis to tabulate where problems with vendors’ compliance lie and, importantly, how each vendor’s compliance and transparency compare with that of the others in the study; we find that for each question asked about the policies, there are varied results. There are areas such as ‘unintended processing’ – when a VA listens and processes data without being asked to – where all four vendors fail.
- Using this information, we develop Privextractor: a web-based software tool that enables users to not only understand the privacy issues that surround the use of VAs but to see, simply and clearly, how the VA that they use compares with others on the market. We see Privextractor primarily as a decision-making tool for use when selecting a VA; with more development of the forensic capabilities of the dashboard Privextractor might become a companion for use throughout the time that the user owns their VA. As far as we are aware, this is the first study to develop such a dashboard.

With this information, as well as the tool Privextractor, we enable users to see that the protection of law such as GDPR does not necessarily bestow the protections that might be expected. Users might not be able to use Privextractor to see, for example, how a vendor is processing their information, and demonstrate that the vendor is not prepared to disclose this information which is a privacy concern. There is much scope for improved law and greater enforcement of that law. It is important that users are able to understand how their personal and private data is being manipulated and, as such, vendors need to improve both compliance and clarity.

1.4 Organisation

The remainder of this paper is organised as follows: Section 2 outlines the methodology for the research and practical elements of the project. Section 3 expands on the methodology by describing exactly how the research and experiments were carried out. The results are shown in Section 4; Section 5 concludes the paper and answers the research questions.

2 METHODOLOGY

To meet the goals of answering the research questions posed in Section 1, the following methodology is proposed. The end solution takes the form of a prototype, proof-of-concept web application called Privextractor. Privextractor – a novel user awareness dashboard – is presented as a web application built using a standard Microsoft technology stack: .NET Core framework using a model-view-controller (MVC) design pattern.

2.1 Comparison Matrix

To present to the user a clear, unbiased picture of the VA vendors' data practices and their compliance with data law, a comparison matrix was developed for which a series of subject areas was chosen. For each subject area, several questions were devised, decomposing the subject area into smaller specific areas of interest; the answers can not only inform a user of the compliance of their chosen device vendor, but also compare that device with others on the market such that the user gains an awareness when selecting a VA.

These subject areas were motivated by a TechDispatch article published by the office of the European Union's Data Protection Supervisor [40]. Whilst not reflective of official data policy, the article outlines areas of privacy concern specifically pertaining to VAs; it is these areas which are specific to the voice interface of the VA that are of particular interest. The individual questions' levels – three for each, denoting a level of compliance – were devised during preliminary research into the vendors' privacy policies. This research enabled us to find the level for each, where one vendor might be transparent and give plenty of compliance information (the 'good' level) and another might give little information (the 'poor' level).

Each of the four vendors' privacy policies, legal notices, and any advertising or cookie-specific disclaimers were examined in detail to gain an insight into how transparently they are written and how much relevant detail is supplied. The questions are then tested against each vendor's policies and the answers are collected and written as objectively as possible to aid further comparison between the vendors.

The answers were compared and each of the four VAs was assessed to give an immediate visual indication of a) how the user's chosen VA complies with data protection law, and b) how the user's choice of VA compares – in terms of vendor transparency – with the other three. The subject areas are as follows:

- **Transparency** – data controllers and processors, types of data processed, purposes of the processing, specific processing of biometric data
- **Consent** – decisions made on processing the data of a specific individual
- **Children** – distinguishing adults from children, age verification, parental responsibility and consent, parental controls
- **Unrequested Processing** – wake word confusion, deliberate wake word tampering
- **Data Repurposing** – data profiling and purposes of profiling, transfer of data to third parties
- **Data Retention** – length of time for which data is kept, user’s ability to control and delete (all or some) data, the ‘Right to be Forgotten’
- **Security** – access control (account and device/app), indications of security technologies employed by the provider to ensure protection of the user’s privacy
- **Government Surveillance** – handling of access requests from law enforcement and government

Each area considers both how the device and/or application operates. Additionally, each subject area considers the privacy policy information for each vendor and presents a clear picture to the user showing compliance (or otherwise) with GDPR. As GDPR and the DPA are almost identical for the purposes of this work, we will focus only on GDPR for the sake of simplicity. Any indication that the vendors had considered the ICO’s age-appropriate code of conduct (where appropriate) was also taken into account.

As well as an immediate visual indication of the user’s chosen VA both in the context of the questions asked of it and the corresponding performance of the other VAs, accompanying information is provided to the user placing the results in context with simple explanations and links to the appropriate data protection law. It is important to note here that a three-stage traffic light approach might seem odd at first glance to someone who practises law; strictly speaking, a legal requirement is either met or it is not. However, in the case of GDPR and the DPA, it will be seen that a law may be interpreted in different ways; this is particularly true if the law is insufficiently explicit in its requirements. With this comparison matrix, we are attempting to demonstrate to the user each vendor’s understanding of the law; the three-stage compliance levels indicate where each vendor is explicit themselves in how they adhere to data and privacy law, and where they might fall down in not imparting sufficient information to the user in their privacy and legal documents.

Privextractor’s user interface allows the user to select ‘their’ VA and always presents the matrix of information relative to that selection.

2.2 Forensic Recovery

In addition to the comparison matrix, we undertook a forensic investigation of Amazon’s Alexa with interesting results. It was found that a user’s data could be recovered – using a session token found in the user’s web browser – from Amazon’s cloud service API. A range of artefacts was found including account information, and recordings of the user’s voice request and their associated text translation made by Amazon’s natural language processing. Alexa’s software-generated ‘replies’ were also found in text form. Similar artefacts were recoverable from Microsoft’s Cortana.

3 PRIVEXTRACTOR DESIGN

This section introduces the methodology used to recover forensic artefacts from two of the VA vendor’s ecosystems; we also introduce the design of the matrix of questions for vendors that forms the basis of Privextractor.

3.1 Forensic Recovery

Preliminary testing revealed that Amazon exposes an API via which user data associated with Alexa interactions can be read. After authenticating to the Alexa web client, the API can be queried and results returned in formatted JSON from the Alexa account. To illustrate, Figure 1 shows an Alexa ‘reminder’ being served to the browser from Amazon’s API.

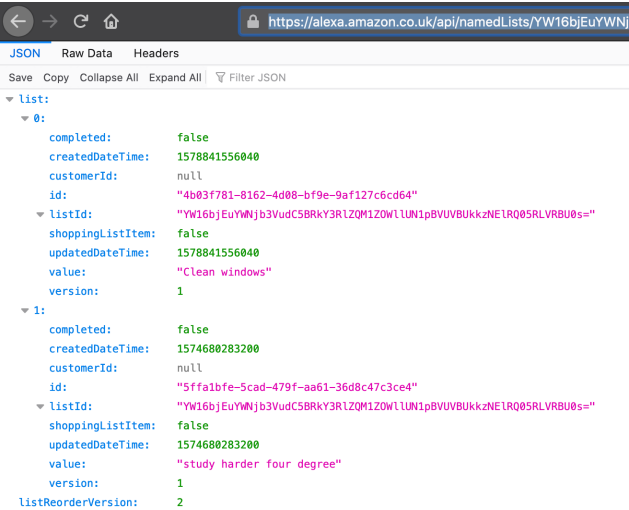


Figure 1: an Alexa ‘reminder’ in JSON format

Microsoft’s Cortana has shown to offer the same facility. Figure 2 shows that, after logging into Microsoft’s Privacy Dashboard, an API can be queried to return data associated with Cortana interactions.

```

{
  "items": [
    {
      "displayText": "What is the weather?",
      "application": "Cortana",
      "deviceType": "Phone",
      "audioId": "29413984EB6D4D5AB28A71DA71878C0F",
      "timestamp": "2020-07-24T16:13:04.4452384+00:00",
      "id": "H4sIAAAAAEACXKMQ6CQBAF0Lv8VjB_Zwd2mG5h9QTGnoAFsSAhdc87o7I",
      "cardType": "cardType_voice",
      "sources": []
    },
    {
      "displayText": "Who is prime minister of the United Kingdom?",
      "application": "Cortana",
      "deviceType": "Phone",
      "audioId": "B92FB924D8C64338AD8A3A4BF0B0851",
      "timestamp": "2020-07-24T15:21:08.9535859+00:00",
      "id": "H4sIAAAAAEACXKMQ6CQBAF0Lv8VjF_Z320drpdCCcg9EYtiIWJoSPcXRI",
      "cardType": "cardType_voice",
      "sources": []
    },
    {
      "displayText": "What is the rweather?",
      "application": "Cortana",
      "deviceType": "Phone",
      "audioId": "73D4297FDD0745AD98715E6B4BDC9A61",
      "timestamp": "2020-07-24T15:20:43.2040061+00:00",
      "id": "H4sIAAAAAEACXKMQ6CQBAF0Lv8VjF_h9kddzpg9ATG3ggFsTAhdIs7Y-I",
      "cardType": "cardType_voice",
      "sources": []
    }
  ],
  "state": null
}

```

Figure 2: Cortana interactions stored by Microsoft and returned in JSON format

Google offers a web-based privacy tool which, upon inspection, retrieves data from an undocumented API to display information pertaining to user interactions with the VA. However, this API was not able to run independently of Google's host code. Apple's Siri has no privacy tool or other client available on the web; all user actions pertaining to Siri's data must take place within native iOS or iPad OS tools.

Amazon, Microsoft, and Google all offer web-based privacy dashboards. Using browser-based developer tools, it could be seen that asynchronous XMLHttpRequest (XHR) requests were being made to the vendors' servers – for example, Amazon's Alexa dashboard was making requests to <https://alexa.amazon.co.uk/api/notifications>. By examining the header for this request, it was possible to see the information that was passed to Amazon's server including, crucially, a session token used to authenticate the API. Similarly, examining the response payload revealed that information was being returned by Amazon's servers. A similar method revealed how Microsoft's privacy dashboard was communicating with its cloud backend, and the information that was being transmitted and received.

Using some code written in Microsoft C#.NET, as part of the PrivExtractor system, it was possible to replicate these requests for both Alexa and Cortana. A valid session token, acquired from the browser whilst logged into Alexa or Microsoft Office365 (in the case of Cortana), was required in order to authenticate. Once an authenticated connection to the APIs could be made, data could be retrieved in JavaScript Object Notation (JSON) format and rendered to the browser.

The data that was recovered from both Alexa and Cortana was real-world data –the information and voice recordings accessible via the API were a result of voice interactions made by the author and, in the case of Alexa, the author's wife using the same device. At this time, the forensics experiments should be considered as a lab study: despite integrating this

facility into PrivExtractor's user dashboard, it is first necessary to obtain a valid session token which is not an everyday user task.

3.2 Comparison Matrix

The eight sections for the comparison matrix to be displayed in Privextractor are described here in further detail. Each section is designed to cover an aspect of GDPR and the UK's DPA; Section 3 (Children) makes additional reference to the ICO's age-appropriate design code of conduct, designed to account for those users of information services aged between GDPR's default of 16 years and the DPA's implementation that defines an adult as 13 years or over.

Each section of the comparison matrix is intended to cover a specific topic, as per the controller's responsibilities laid down in Article 24. Having studied and understood each of the four vendors' privacy policies, legal notices, and any advertising or cookie-specific disclaimers, that insight is used to devise a series of questions. These questions cover a broad range of areas within each specific topic and are intended to indicate the vendors' level of compliance with data law.

This is a qualitative assessment. To provide a measure of visual information, however, each question is assessed in terms of the information provided by the vendor and how this appears to meet the requirements of data legislation. Each question will be given a scale of three possible scores, of 'good', 'average', or 'poor'. The definitions for each will be clearly signed.

The scores are not intended to be an immediate indication of quality when taken in isolation, as might be awarded to a product in a consumer magazine review. However, taken together, the scores show broadly how each vendor is committing to data protection law and, crucially, indicate to an end user how their choice of VA performs when compared with others on the market. The results are tabulated in much the same way that commercial risks are evaluated using a matrix in ISO27001 [41]; the final tables give a clear, colour-coded indication of performance both in isolation and in the context of other VA devices.

3.2.1 Transparency

Question 1: Is it clearly stated who the data controllers/processors are?

- (Good) Yes – name and address
- (Average) Yes – name only
- (Poor) Not stated

GDPR makes specific definitions of 'controller' and 'processor'. This question asks if the vendors specifically outline which parts of their businesses are responsible for each role and if any detail is given.

Question 2: Are the types of data processed – such as a user's name or location data - clearly listed?

- (Good) Yes – examples are given covering GDPR
- (Average) Yes – generic classifications only, incomplete coverage of types stated in GDPR
- (Poor) No – data types not listed, even in generic form

GDPR gives a list of examples of ‘personal data’ that might be taken from an individual during the use of their products. Question 2 asks how specific the vendors are when giving examples of those types of data that might be collected from a user.

Question 3: Are the purposes of processing clearly listed?

- (Good) Yes – examples given covering GDPR
- (Average) Yes – generic classifications only, incomplete coverage of types stated in GDPR
- (Poor) No purposes given, even in generic form

GDPR gives clear examples of what it considers to be the ‘processing’ of data; these range from the simple act of collecting the data in the first instance, to disposal at the end of the process. Question 3 asks how specific the vendors are when outlining the processing purposes.

Question 4: Is any processing of biometric data clearly explained?

- (Good) Yes – examples given covering GDPR
- (Average) Yes – generic examples only, incomplete coverage of types stated in GDPR
- (Poor) No information about biometric data processing is given

GDPR has a definition of what constitutes ‘biometric data’. Voice recordings might not appear as categorically ‘biometric’ as, say, a fingerprint or retinal scan; however, each of the VA vendors does engage in some form of fingerprinting – identifying a person using their data – to personalise the user experience upon recognising their voice. This voice fingerprinting process has a significant precedent: when the UK’s governmental tax collection department – Her Majesty’s Revenue and Customs (HMRC) – adopted voice authentication in 2017, complaints were made by industry watchdogs due to the lack of transparency from HMRC [42].

3.2.2 Consent

Question 1: Does the device feature a mechanism whereby it processes the data of only a specific individual?

- (Good) Yes – data processing limited to a single user at the device level
- (Average) Only for specific features, or for personalisation
- (Poor) No mechanism offered – the device will process the data of any user who interacts with it

For the controller to process data, consent must be obtained from the user. This is important enough for GDPR to define what it means by ‘consent’. The ‘data subject’ is the one it is assumed has given consent; another person using the same VA might not have done. Currently, none of the four VAs has the facility to perform voice identification without sending the recording to the vendor’s cloud service, at which point the transmission of the data as well as the analysis at the vendor’s end is considered ‘processing’ by GDPR. It is possible for this to happen without consent having been given. However, as VAs become more capable at the device level, Privextractor will be updated and the results of this question might change.

3.2.3 Children

This section of the matrix introduces the ICO’s age-appropriate design code. This code is not enshrined in UK law, rather it sets standards and explains how UK GDPR ‘*applies in the context of children using digital services*’ [25]. Whilst GDPR

considers an ‘adult’ to be anyone over the age of 16 years, the DPA lowers this to 13. The ICO’s code helps bridge this gap of three years with advice to providers of digital services whose services might be either aimed at children or whose services might reasonably be accessed by a child.

Question 1: Does the provider distinguish between adults and children as users?

- (Good) Yes – with age explicitly stated
- (Average) Yes – no age stated
- (Poor) No distinction made based on the user’s age

As a baseline for Question 2, this question asks if the vendor states, in their privacy policies, what they consider to be the age of an adult as distinct from a child?

Question 2: If applicable, what form does the age verification mechanism take?

- (Good) External verification using endpoints not easily obtainable by children (credit card)
- (Average) Basic input of age, with external verification using endpoints easily obtainable by children (email, SMS)
- (Poor) External verification only as a means of two-factor authentication, age not considered

Background: GDPR requires that controllers ‘shall make reasonable efforts’ to verify the age of the primary user during initial setup, or that consent is given by the responsible adult. The ICO’s Age-appropriate design code goes further and suggests some mechanisms by which this might be done, from simple self-declarations to more complicated credit card checks. However, even the strongest of these verification methods is not without issue; whilst credit card checks are appropriate for children, they pose a problem for those aged between 13 years and 18 years who are considered adults by the DPA but cannot – in the UK – legally hold a credit card with which to verify their age. For reference, the section in the DPA which deviates from GDPR for the UK is shown in Table 1.

Table 1: excerpts from DPA Section 9

‘Child’s consent in relation to information society services’: In Article 8(1) of the GDPR (conditions applicable to child’s consent in relation to information society services)

(a) references to “16 years” are to be read as references to “13 years”, and

(b) the reference to “information society services” does not include preventive or counselling services.

Question 3: Is there a way of ensuring the person with parental responsibility has provided consent for a child’s interaction with the device?

- (Good) Yes – by full authorisation
- (Average) Yes – by optional ‘parental’ mechanisms
- (Poor) No mechanism present for giving parental consent

This question asks if it is possible that when a child is using the device after its initial setup, a parent can be assumed to have given consent. There are ways in which this might happen, for example by the use of optional parental controls offering the parent or guardian the ability to limit when the child uses the device or service.

Question 4: Are there any parental controls?

- (Good) Yes – fine-grained control on all devices
- (Average) Yes – some control, or only on certain devices
- (Poor) No parental controls present

The ICO's age-appropriate design code offers useful insight into parental controls that can be "*used to support parents in protecting and promoting the best interests of their child*" [25]. Does the vendor offer any controls and, if so, do they operate across all devices on which their VA application might reasonably be expected to be used?

Question 5: Are the parental controls made available with good accessibility for users?

- (Good) Yes – clear instructions signposted in online support
- (Average) Yes – but the information is difficult to find
- (Poor) No – there is no information given regarding the controls

Parental controls are of little utility if they are hard to operate, or information explaining how they work is difficult to find. For the purposes of this question, 'difficult to find' means the information is not clearly available from the vendor's online support. There is a further nuance – the child whose access is being controlled has, under GDPR's edict that personal data shall be "*processed lawfully, fairly and in a transparent manner in relation to the data subject*", a right to know how the controls are affecting their use of the service. The ICO's code suggests that, for children aged between 13-15 years, the vendor's information should also clearly explain this.

3.2.4 Unrequested Processing

Question 1: Is there evidence that a mistake could be made, confusing a spoken expression for the VA wake word?

Question 2: Is there evidence that VAs mishear their wake words, leading to accidental recordings?

Background: VAs are activated by a wake word, after which an indication is given that it is ready for the user to interact with it. To know if the wake word has been spoken by the user, the device needs to be constantly aware of the sounds being made near it. A VA should not be recording or sending any information to its cloud server, however, until the wake word has been spoken. Each device has a wake word; Alexa offers the facility to change the word from a predefined selection. The words are 'Alexa', 'Echo', 'Computer', 'Ziggy', or 'Amazon' for Alexa devices; 'Hey Siri' for Apple devices; 'Hey Google' for Google devices; 'Hey Cortana' for Microsoft devices.

Should the device mishear the wake word, it might activate and start sending audio to the cloud for recording without the user's knowledge – a clear privacy breach. VAs use audible alerts and visual indicator lights to mitigate the chance of this happening without the user's knowledge, but these only alert the user to the fact that a recording is being made, they do

not prevent it. The results of these questions are not presented in a chart but simply answered ‘yes’ or ‘no’ with evidence, if applicable, to assert the answer.

3.2.5 Data Repurposing

According to a report made by Reuters Practical Law, ‘Big Data’ relies on three things: aggregation (size – vast volumes, shape – text, sound); analysis (datasets are analysed in real time); and increasing value (enhancing competitiveness and efficiency) [43]. Data is not in short supply for these companies and their VAs are adding to it; these questions ask if the VA companies explicitly state how they repurpose or share this information.

Question 1: Are the purposes of any data profiling explicitly stated?

- (Good) Yes – examples given covering those explicitly stated in GDPR
- (Average) Yes – generic classifications only, and/or incomplete coverage of purposes
- (Poor) No examples or purposes of data profiling given

As distinct from what GDPR calls ‘processing’, which can be as simple as the collection of the data in the first instance, profiling refers to the manipulation and mining of the data to infer the characteristics of the user. A common use for this is targeting advertising, where the user’s interests have been built into a profile that matches that of a seller’s target – someone who may be susceptible to buying the seller’s product.

Question 2: Is the user’s data – according to the vendor’s policies – shared with other entities outside of the organisation?

- (Good) Yes – with explanations of what is shared, why, and with whom
- (Average) Yes – no explanation given
- (Poor) No – the user’s data is not shared outside the organisation

Recital 6 of GDPR highlights sharing of data as a growing area of concern, and one of the drivers in the introduction of the regulation.

3.2.6 Data Retention

Question 1: Can users find out how long data will be stored?

- (Good) Yes – with specified timescales
- (Average) Yes – without specified timescales but within parameters of certain events
- (Poor) No – users are unable to find out how long their data will be stored for

GDPR is specific about how long user data should be kept for. Exceptions are made for cases where users’ data is processed for archiving purposes in the public interest, for scientific or historical research purposes, or for statistical purposes. As it is unlikely that Amazon, Apple, Google, or Microsoft are engaged in these activities, they are obliged to keep it for no longer than they need to process it.

Question 2: Is it possible for a user to delete their voice recordings?

- (Good) Yes – clearly signposted in online support

- (Average) Yes – not clearly indicated in help guides
- (Poor) No – users are unable to delete their own voice recordings

When a user interacts with a VA, a recording is made of their voice and sent to the vendor's server for processing. These recordings are kept, in the form of an audio file, alongside the user's account. It is important to note that it is not possible to prove that a recording has been deleted – with this question, we are taking the vendor at their word. Even if a recording appears to have been deleted, the vendor may have simply removed it from visibility of the user.

Question 3: Does the delete function remove all data (transcriptions) or just voice?

- (Good) Yes – all data
- (Average) Voice data only
- (Poor) Some voice data cannot be deleted

Alongside the audio recordings of the user's voice interaction, a text translation is made by the vendor's speech recognition software. However, as is shown in the 'Repurposing of Data' section, the user's information is not just used for responding to queries. The data is used for advert profiling, 'personalisation', and any manner of other purposes; as long as the user still consents to these practices, the providers are entitled to store the data.

Question 4: Does the provider offer 'The Right to be Forgotten'?

- (Good) Yes – clearly signposted with a selection of contact routes (verbal, writing)
- (Average) Yes – limited means of request
- (Poor) No – the provider does not offer the right to be forgotten

GDPR and the DPA offer what is called 'The Right to be Forgotten' which obliges the controller, when requested by the user, to erase the user's personal data '*without undue delay*'. This can be triggered in several ways, for example where the processing of the data is found to be unlawful or if there is a national or EU legal obligation to do so. Where VAs are concerned, the important reason is when '*The data subject withdraws their consent and the controller has no other legitimate ground for the processing of the data.*' In other words, when the user has decided that they no longer want the provider to keep their data and withdraws their permission for the provider to do so. Users are entitled to be able to make the request (withdraw their consent) verbally or in writing.

The right to be forgotten made news when, in 2019, Google fought the EU Court of Justice and won a landmark ruling against the French privacy regulator Commission Nationale de l'Informatique et des Libertés (CNIL). The outcome of the case was that Google only had to oblige the user's right to be forgotten in EU countries and not globally. Google argued that they didn't wish to see totalitarian governments forcing their political will on their populations by removing and therefore skewing search results in their favour [44].

3.2.7 Security

Where personal data is concerned, security is paramount in order to ensure the user's data remains private. VA devices themselves have been the target of malicious attacks, as reported by various news agencies [45].

Question 1: Is there any access control (authentication, authorisation) to the provider account?

- (Good) Yes – credentials and 2FA
- (Average) Yes – credentials only
- (Poor) No access control in place

As it is necessary to create an account with each of the vendors, the security of that account is important. Access to the account could give an intruder personal and private data; moreover, anyone with control of that account could use it to impersonate the original user causing financial and reputational loss – some VAs allow the user to make purchases by voice.

Question 2: Is there any access control to the VA device or app?

- (Good) Yes – the vendor’s VA is protected on all devices
- (Average) Yes – the vendor’s VA is only controlled on some compatible devices
- (Poor) No access control in place

Question 2 deals with the security on the VA device or application itself as opposed to the security protecting the user’s account. With the proliferation of ways in which one single VA – i.e. Google Assistant – can be used, on smartphones, tablets, and smart speakers, the methods in which the VA may be secured vary. A smart speaker may not have any inbuilt security, allowing it to be used by anyone in its vicinity; however, a VA used on a smartphone may be protected by the phone’s security, in the form of a PIN code or a fingerprint scan. GDPR is unclear on this definition – there is little suggestion of how the controller might be required to implement any security on its endpoint software which has access to its cloud servers.

Question 3: Does the provider indicate that encryption is used for the protection of data in transmission or when stored?

- (Good) Yes – examples of technologies given
- (Average) Yes – no specific detail provided
- (Poor) No information given regarding security in transit or at rest

GDPR’s main focus – when discussing security – is the technologies the vendor uses to protect data in storage and data in transit to ensure the information remains private. Question 3 asks if the vendors are upfront and give examples of the security methods they use, and how specifically those measures are communicated to the user. Problems are further compounded – and GDPR does make specific mention of this – when employees at the vendor are given access to voice recordings [46].

3.2.8 Government Surveillance

Question 1: Is the user informed if the vendor discloses information when an access request is made by law enforcement or government agencies?

- (Good) Yes – clearly states the user will be informed, and how
- (Average) Yes – no detail given
- (Poor) No – the user is not informed

In 2016, the UK Government's then Home Secretary Theresa May introduced the Investigatory Powers Act (IP Act). This act gave UK intelligence agencies (including MI5), and law enforcement, new powers to carry out interception of communications and to collect communications data in bulk [47]. The London School of Economics believed at the time that the IP Act could conflict with GDPR [48]. It will be of interest to see the vendors' practices in this regard.

4 RESULTS AND TESTING

In this section, we carry out the research required to answer the questions that form the matrix designed in Section 3: Privextractor Design. The answers to the questions, obtained from the companies' privacy policies and legal statements, are tabulated and shown as part of the user interface in the final software application. Additionally, we show how the findings from Section 3.1: Forensic Recovery are incorporated into Privextractor's dashboard.

4.1 Comparison Matrix

A sample of the matrices discussed in this section can be seen here as displayed in Privextractor's user interface.

4.1.1 Tabulated Results

The following sections outline the findings from examining the vendors' privacy policies, and answering the questions posited in each area of the matrix designed in Section 3: Privextractor Design. In each table, a colour scheme is used where green (happy face) = good, yellow (indifferent face) = average, and red (sad face) = poor. This colour scheme gives an immediate visual indication of the standard of each VA vendor's policies when asked a specific question. The scores are intended as a comparison of the vendor's privacy practices; should two VAs achieve the same score, the user can see that choosing one of the VAs means that there is no advantage in either selection.

Firstly, Figure 3 demonstrates the user selecting their VA – in this case, Google Assistant.




Figure 3: user VA selection in Privextractor

In Figure 4 it can be seen that the user has selected Google Assistant and the score for that assistant is now highlighted. Below can be seen the matrix for the first question in the 'Transparency' section.

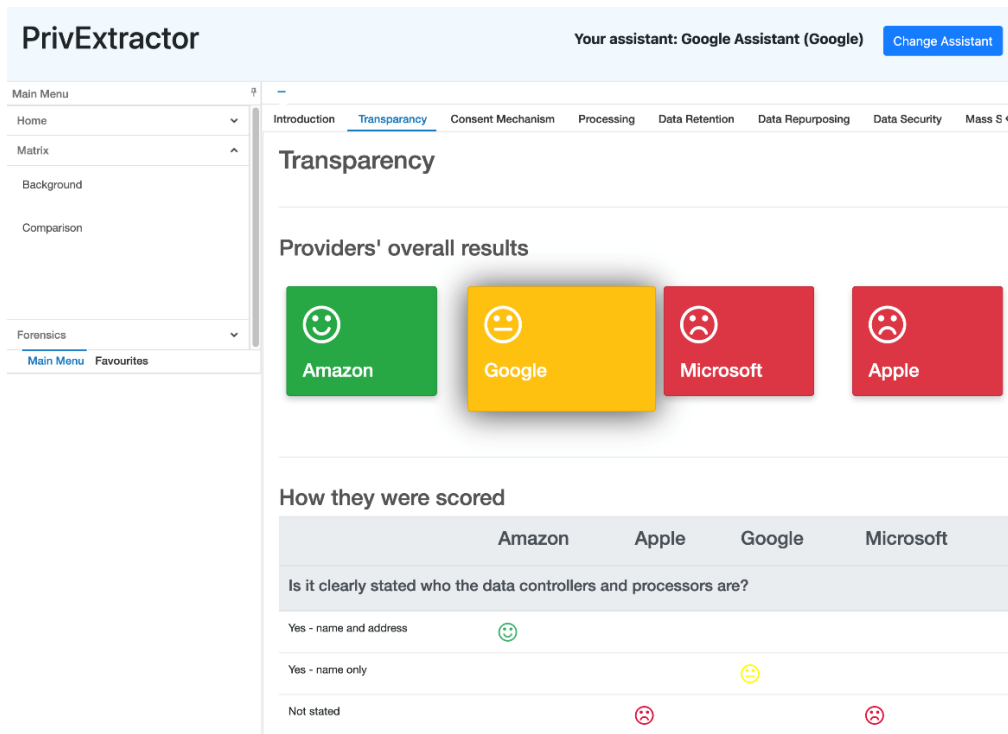


Figure 4: Transparency comparison with Google Assistant selected

Finally, in Figure 5, an expanding box has revealed the information used to populate the matrix for the first question.

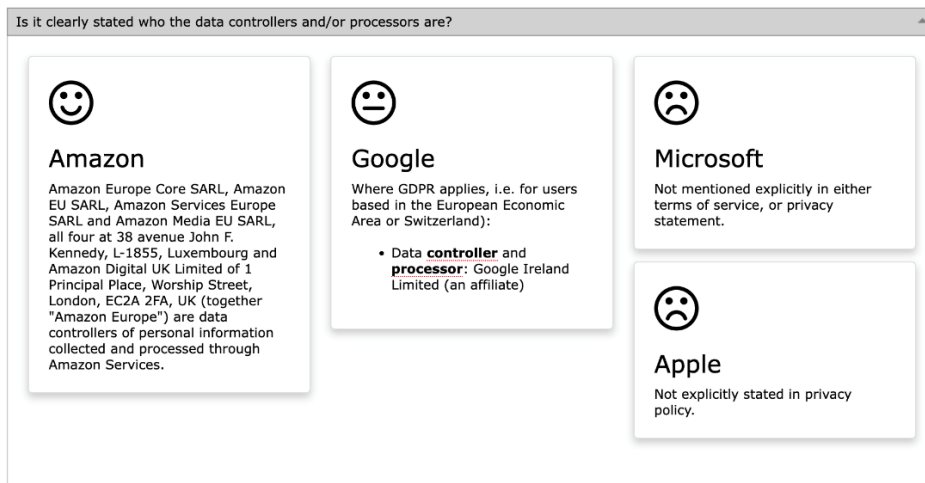


Figure 5: Transparency comparison with Google Assistant selected

4.1.2 Transparency

Table 2 shows the results of the first section of the matrix in which basic definitions were examined in each vendor's privacy policies to compare how open each was about processors and processing, and the types of data that are collected by each.

Table 2: a comparison of vendors' transparency

	Amazon	Apple	Google	Microsoft
Is it clearly stated who the data controllers/processors are?				
Yes – name and address	😊			
Yes – name only			😐	
Not stated		😞		😞
Are the types of data processed – such as a user's name or location data - clearly listed?				
Yes – examples are given covering GDPR		😊		😊
Yes – generic classifications only, incomplete coverage of types stated in GDPR	😐		😐	
No – data types not listed, even in generic form				
Are the purposes of processing clearly listed?				
Yes – examples given covering those explicitly stated in GDPR				
Yes – generic classifications only, incomplete coverage of types stated in GDPR	😐	😐	😐	😐
No purposes given, even in generic form				
Is any processing of biometric data (voice) clearly explained?				
Yes – examples given covering GDPR	😊			
Yes – generic examples only, incomplete coverage of types stated in GDPR				😐
No information about biometric data processing is given		😞	😞	

4.1.3 Consent

Here we looked at the ability of the VA to respond only to its original user. As per Table 3, only Microsoft's Cortana had this feature; however, it should be noted that Microsoft's disclaimer stated Cortana will 'try' to respond only to the user whose voice it has been trained to recognise.

Google offers 'Family Link' with which its VA can be trained to respond to a child in the family as part of a wider set of parental controls. Amazon's Alexa and Apple's Siri offer voice training that will personalise their responses to the user whose voice they recognise; they will not, though, prevent a stranger from conversing with the system.

Table 3: a comparison of VA consent mechanisms

	Amazon	Apple	Google	Microsoft
Does the device feature a mechanism whereby it processes the data of only a specific individual?				
Yes – data processing limited to a single user at the device level				😊
Only for specific features, or for personalisation	😞	😞	😞	
No mechanism offered – the device will process data of any user who interacts with it				

4.1.4 Children

This section looks at the vendors' privacy policies compliance not just with GDPR, but with the ICO's Age-appropriate code of conduct. This code was written to cater for a child's use of information services in general, but specifically, those users who are still considered children by GDPR (aged 16 or under) but not by the DPA (13 or under). The ICO make recommendations about age verification, consent and parental controls. The results can be seen in Table 4.

Table 4: a comparison of VA vendors' practices as pertain to child users

	Amazon	Apple	Google	Microsoft
Does the provider distinguish between adults and children as users?				
Yes – with age explicitly stated	😊	😊	😊	
Yes – no age stated				😞
No distinction was made on the user's age				
If applicable, what form does the age verification mechanism take?				
External verification using endpoints not easily obtainable by children (credit card)				
Basic input of age, with external verification using endpoints easily obtainable by children (email, SMS)		😞	😞	😞
External verification only as a means of two-factor authentication, age not considered	😞			
Is there a way of ensuring the person with parental responsibility has provided consent for a child's interaction with the device?				
Yes – by full authorisation			😊	
Yes – by optional 'parental' mechanisms	😞	😞		😞
No mechanism present for giving parental consent				
Are there any parental controls?				
Yes – fine-grained control on all devices		😊	😊	
Yes – some control, or only on certain devices	😞			😞

No parental controls present				
Are the parental controls made available with good accessibility for users?				
Yes – clear instructions signposted in online support		😊	😊	😊
Yes – but the information is difficult to find	😐			
No – there is no information given regarding the controls				

4.1.5 Unrequested Processing

All four VAs failed this test - research has shown evidence that all four can mishear the wake word, and record without the user's consent and knowledge – for example when a user says something that the VA thinks is its wake word, or when the VA hears something on the television or radio that it mistakes for its wake word [49].

Only Amazon allows the wake word to be changed – and then, the choice is limited to one of four options. It is worth noting that this cannot be done on every device running Amazon's VA, Alexa [50].

4.1.6 Data Repurposing

As explained in Section 1 - Introduction, the repurposing of data – also known as 'mining' – is a profitable business for the vendors. The vendors' comparative transparency can be seen in Table 5.

Table 5: a comparison of VA vendors' data profiling policies

	Amazon	Apple	Google	Microsoft
Are the purposes of any data profiling explicitly stated?				
Yes – examples given covering those explicitly stated in GDPR		😊	😊	😊
Yes – generic classifications only, and/or incomplete coverage of purposes	😐			
No examples or purposes of data profiling given				
Is the user's data – according to the vendor's policies – shared with other entities outside of the organisation?				
Yes – with explanations of what is shared, why, and with whom		😊		
Yes – no explanation was given	😐		😐	😐
No – the user's data is not shared outside the organisation				

4.1.7 Data Retention

Data retention – specifically for how long VA vendors keep a user's information – is a key tenet of GDPR. The vendors all showed broad compliance, apart from Microsoft (refer to Table 6). GDPR makes a provision for the 'right to be forgotten' which was mentioned in three of the vendors' policies. GDPR makes no provision, however, for allowing the interim deletion of data by the user. All but Amazon failed here, and even they allow only the deletion of voice recordings.

Table 6: a comparison of vendors' practices regarding data retention





	Amazon	Apple	Google	Microsoft
Can users find out how long data will be stored?				
Yes – with specified timescales			😊	
Yes – without specified timescales but within parameters of certain events	😐	😐		
No – users are unable to find out how long their data will be stored for				😞
Is it possible for a user to delete voice data?				
Yes – clearly signposted in online support	😊	😊	😊	😊
Yes – not clearly indicated in help guides				
No – users are unable to delete their own voice recordings				
Does the delete function remove all data (transcriptions) or just voice?				
Yes – all data				
Voice data only	😐			
Some voice data cannot be deleted		😞	😞	😞
Does the provider offer 'The Right to be Forgotten'?				
Yes – clearly signposted with a selection of contact routes (verbal, writing)	😊			
Yes – limited means of request		😐	😐	
No – the provider does not offer the right to be forgotten				😞

740 4.1.8 Data Security

741 The VA vendors' compliance with GDPR's security requirements is examined and the results shown in Table 7.

742 Table 7: a comparison of vendor and application security
743





	Amazon	Apple	Google	Microsoft
Is there any access control (authentication, authorisation) to the provider account?				
Yes – credentials and 2FA	😊	😊	😊	😊
Yes – credentials only				
No access control in place				
Is there any access control to the VA device or app?				
Yes – the vendor's VA is protected on all devices				
Yes – the vendor's VA is only controlled on some compatible devices	😐	😐	😐	😐
No access control in place				

Does the provider indicate that security is used for the protection of data in transmission or when stored?				
Yes – examples of technologies given				
Yes – no specific detail provided				
No information was given regarding security in transit or at rest				

4.1.9 Government Surveillance

As seen in Table 8, Microsoft did not mention in their privacy statement whether or not they inform users – in this case, it is assumed they do not. Amazon expressly stated that they do not inform the user.

Table 8: a comparison of the vendors' practices in dealing with access requests

	Amazon	Apple	Google	Microsoft
Is the user informed if the vendor discloses information when an access request is made by law enforcement or government agencies?				
Yes – clearly states the user will be informed, and how				
Yes – no detail was given				
No – the user is not informed				

5 DISCUSSION

The results obtained during this research raise several questions. The comparison matrix, ostensibly designed to test the VA vendors' compliance with data law, has done just that. Whilst there is room for improvement in specific areas of the vendor's adherence to data law, it has been shown that it is not so much the vendors' compliance that is of concern but the law itself. GDPR has proved to be quite vague in several areas, meaning that its purpose – to protect the user – is failing.

Whilst we feel that criticism of the vague nature of the requirements laid out in GDPR and the DPA is valid, it should be pointed out that these regulations are, by necessity, designed to handle a large number of divergent cases of which user data exchanged with a VA is just one. This does not negate any criticism of under-regulation inherent in current data law, and it is clear that GDPR and DPA must be regularly updated. Laws are not concrete and are open to interpretation, but they must be considered in a way such that they provide a solid foundation for protecting the user.

Whilst complying with GDPR, the vendors are acquiring large amounts of data and are not specifically informing the user what they are doing with it. Despite declaring that they do not 'sell' data, the vendors are exchanging information for money via advertising platforms. GDPR could improve in this area and require the vendors to explicitly state how and when this happens, and when they profit, which would improve on the generic caveats given currently in privacy statements. The understanding of terms by the vendors must be as precise as possible – particularly here, where the terms are applied in a specific case. It must not be possible for the companies to take a divergent line in their own legal terms and conditions.

Advertising is not the only issue. It has been seen that the UK government has previously collected the communications and social media data of its citizens [51]; should they come into possession of the information collected by VAs, this could be considered a worrying breach of privacy. The regulation allows ambiguity in the vendors' outlining how and when information is shared with law enforcement and governments; two of the vendors openly admit that the user will not even be informed when their personal data is shared with a government agency.

Moreover, GDPR is very specific about requiring user consent without offering any concrete guidance on how this might be obtained by VAs. Again, regulation that deals with divergent cases – as here, where many devices are covered – must be neutral to the technology. However, if consent cannot – for any reason – be effectively given, then users, in particular children, are inevitably going to have their data processed without their consent and having no knowledge of the privacy policies governing their data.

5.1 Privextractor

In order to convey this information to the user of a VA, we developed Privextractor – a web-based dashboard comparison tool. Privextractor contains the information outlined in Section 4.1: Comparison Matrix and – via a mechanism whereby the user can choose their choice of VA – offers an easy reference comparison for the user to decide how the vendor of their VA is complying with data law. This information can help empower the individual to learn how their data is treated, and when and with whom it is shared.

We envisage this tool to be used in two ways: firstly, as a reference point for a user to select a VA and, secondly, as a tool for the user to reference throughout the time that the user interacts with their VA. We have seen that, whilst vendors are largely compliant with data law, the law itself is not specific enough to enforce transparency on the part of vendors such that users have a full and honest picture of what is happening to their data. A tool that can help redress this balance will enable the user to interact with their VA in greater confidence. Privextractor is, we believe, the first of its kind to offer this facility. We have seen from Section 1.1: Related Work that there are studies dedicated to user perceptions of VA privacy, and that those users are concerned; Privextractor, we hope, will help to address those concerns when a user chooses a VA.

During the user VA lifetime, a guide to how the user's data is being collected gives a useful overview to the information that the user has shared with the device and, by extension, the vendor. Zibuschka et al. aimed to expose this information in their own dashboard – ENTOURAGE [16]; we feel that the combination of this forensic work, already begun in Privextractor, and the comprehensive overview of the vendor's privacy practices with regard to data law make a useful 'one-stop shop' that can act as a reference point for the VA user as long as they use their device.

Whilst these use cases, we feel, are advantageous to PrivExtractor's target audience – end users – there are also shortcomings of such a system. The four vendors whose VAs are studied here – Apple, Amazon, Microsoft, Google – are large and dynamic organisations whose legal terms and privacy policies are likely to change regularly. Privextractor, in its current state, cannot dynamically accommodate such alterations to this source material in which it is based. Moreover, its forensic capability is – at present – limited to what can be thought of as a laboratory experiment, due to the complications with locating and identifying the necessary security token required to authenticate to the vendors' cloud services.

PrivExtractor’s main reason for existence is to assist users in making informed choices about their use of a VA and, as such, its utility for practitioners such as regulatory bodies and the vendors themselves is, on first inspection, limited. However, a study by Emami-Naeini et al. determined through interviewing users that there are many considerations a user makes when purchasing a VA, amongst which privacy and security concerns were only mentioned by a few participants [52]. Post-purchase, the number of participants reporting security and privacy concerns rose to around half of the total number of interviewees. The authors proposed the use of a label, to be attached to the device at point of sale, containing “...ratings from an independent privacy lab, an independent IT security institute, and Consumer Reports (CR).”

Should a regulatory body start to demand such transparency, a tool such as PrivExtractor could work in tandem with such a system with the labelling offering a useful insight into the VAs security and privacy rating at point of sale, with PrivExtractor as a companion throughout the life of the VA. Projects such as Polisis [18] which aim to automate privacy policy analysis, and Privacy Flag [17] which assess a user’s smartphone apps for privacy risk, could be further systems with which PrivExtractor might work. A painstaking and time-consuming part of this research was the manual analysis of privacy policies, and a sophisticated automated system such as Polisis would be a great asset. Furthermore, VAs – as we have seen – exist not only as standalone devices, but as smartphone applications. Privacy Flag’s important work in determining the privacy risk of smartphone apps could be a very useful system with which PrivExtractor may interact. In addition, we would be happy to share usage statistics and user feedback from the use of PrivExtractor with VA vendors should they wish.

In academic research terms, we saw in Section 1.1: Related Work how Ford et al. attempted to analyse the traffic that is exchanged between Amazon’s Alexa VA and its cloud platform [35]; unable to decrypt the TLS-encrypted traffic itself, the authors had to resort to observing patterns in the quantity and timing of the data that passes across the network. PrivExtractor continues the forensic work that has been carried out in previous studies [36] [37] [38]; our forensic work does not offer anything new, but reinforces existing studies which helps in an understanding of how VAs handle data and what the vendors are storing. We also saw in the literature review how many academics are interested in user perceptions of VAs, and how they view these devices in privacy terms. We see PrivExtractor as a useful tool to help in this area, by addressing one of our goals of redressing the imbalance of understanding between the user and the VA vendors.

5.2 Research Questions

In this section, each of the research questions is addressed individually.

RQ1: If a user of a voice assistant wishes to inform themselves about the extent to which their personal information is harvested by the vendor of their chosen VA, does that vendor clearly and unequivocally state the exact nature of the information that they collect, how securely they keep that data, what they are doing with it, and for how long they keep it?

As can be seen in section 4.1: Comparison Matrix, where the results of this part of the research are described in full, the vendors of virtual assistants (VAs) are largely compliant with data law and any deviations from the strict rule are minimal. It is where GDPR itself becomes less clear that a corresponding lack of clarity is found within the vendors’ privacy policies. For example, Google’s privacy statement appears quite specific in defining what the company does not share – “We don’t

show you personalized ads based on sensitive categories, such as race, religion, sexual orientation, or health” is one example.

As described in Section 1: Introduction, Google makes the majority of its money by brokering online advertising [6]. The user, however reassured that they will not be shown adverts based on their race, for example, might still like to know the following:

- What data *do* you share?
- Exactly when do you share the data?
- With whom?
- Is the sharing for profit?

Amazon, similarly, gives plenty of information in its privacy policy that suggest it complies with GDPR’s requirements; however, there are apparent contradictions. The company made over US\$10 billion in ad revenue in 2019 [53] despite the claim made in their privacy policy that they “*are not in the business of selling our customers’ personal information to others*”.

More clues are given in Amazon’s ‘Interest-Based Ads’ policy in which the company tells the user that they ‘*work with*’ third parties; from this, it is not clear if they mean ‘*share data with*’. These examples are in accordance with GDPR and, as such, neither Amazon nor Google is in breach of the law. The lack of information that is conveyed to the reader, though, suggests that in certain areas GDPR is not providing the overarching protection of the user that it is intended to.

RQ2: Using UK and EU data law as a basis – GDPR, the DPA and the ICO’s age-appropriate code of conduct – can it be demonstrated to a user where the data collection practices of the VA vendor conflict with the law?

All four vendors perform reasonably well when questioned on their data collection practices in terms of what is required to be outlined by GDPR. Some vendors perform better than others in certain areas – they provide more detail, but GDPR as it stands is being adhered to. Despite this, there do remain some areas of concern, chiefly around the way the vendors handle security, authorisation, and consent.

In general, there are two points where security might be a concern in the use of a VA; the first of which is the security of the vendor-hosted service account that it is necessary to create in order to use the VA, and the second is the security of the VA application itself. Both of these must be robust in order to ensure that any information exchanged with the device and the vendor remains private. Transparency regarding account security is handled well by all the vendors, as is the security itself: all vendors offer two-factor authentication, and these accounts are well-protected against a malicious threat actor wishing to gain illicit access to a user’s personal information. There is some opacity in the information provided by the vendors regarding the means of securing cloud data once it is in the possession of the vendor, but some information is given.

Where all four vendors fall is access control to their VA application across devices. There is simply no way, on any of the four vendors’ VAs, of ensuring that – at all times – consent has been given by the person interacting with the client

application. Whereas the individual who initialised the device and created the cloud service account has given consent during the signup process, any further user who interacts with the VA has not. This is less of a problem when the VA is used on a smartphone or tablet, as many of these offer device-level security such as PIN codes or fingerprints to access. Smart speakers, however, do not.

Whilst a concern for any user, lack of consent and authentication mechanism becomes more problematic when the subsequent user is a child. Amazon, for example, clearly states *“If you're under 18, you may use Amazon Services only with the involvement of a parent or guardian.”* The ICO’s view is that children aged between 16 and 17 years *“are still developing cognitively and emotionally and should not be expected to have the same resilience, experience or appreciation of the long term consequences of their online actions as adults may have.”* [25]

All vendors distinguished children from adults in their privacy policies, with only Microsoft failing to state the age of what it considers a child. All four vendors offered parental controls; Google’s controls in particular – ‘Family Link’ – are fine-grained and offer voice fingerprinting as a mechanism of making sure the child is correctly authenticated.

There is no mechanism preventing a child from signing up for an account to use any of the assistants that they could not easily circumvent. Other options available to the vendors, such as making a user confirm their adulthood via a credit-card check, have been the subject of much debate by owners of pornography websites [54] so it is perhaps unfair to expect similar implementation for a service such as a VA that a child may be reasonably expected to use. In the UK it is illegal to obtain a credit card until the age of 18 which – given the age of an ‘adult’ in terms of some of the VA vendors is a maximum of 16 years – presents a further issue.

Other checks recommended by the Information Commissioner’s Office for verifying the age of a user include artificial intelligence or ‘hard identifiers’ such as a passport [25]. These could be considered equally obtrusive, further eroding privacy; relying on mandatory confirmation of an adult account holder would appear to be a better compromise.

The fact remains, however, that VAs gather a lot of data – voice, geolocation – that when taken from a child could present a safeguarding risk [55]. Whilst compliant with data law, simply relying on a self-declaration is insufficient to mitigate this risk and would appear to represent a shortcoming in the law itself.

5.3 Conclusion

Privextractor is a novel proof-of-concept application that is capable of highlighting to a user the strengths and weaknesses of a chosen VA. From a review of the literature this has not previously been reported; there has been little work undertaken on a self-contained user awareness tool specifically targeting virtual assistants. Such a tool could significantly increase VA users’ understanding of the privacy and security issues surrounding the use of an assistant.

The outcomes of the work are interesting – we started the research with an open mind, and did not know what to expect. Two possibilities were that the vendors were a) complying with data law and there was no problem, or b) were not complying with data law causing an obvious legal issue. Curiously, the outcome was, strictly speaking, neither – the vendors are in compliance. However, the more we researched and studied this area, the more it became apparent that data law such as GDPR is not specific enough to allow the user to make an informed choice in the VA market should they be

concerned about privacy. Ultimately, the law has to cover a lot of different cases and cannot be too specific – but we feel a lack of specificity in what are quite tightly-defined areas (“We do not sell your data”) is allowing the VA vendors too much latitude at the user’s expense.

Previous studies have examined user awareness and acceptance factors of VAs [26] [27]. Studies have also been made on the forensic recovery of information from Amazon’s cloud service, work which resulted in a functional web application [37]. However, there have been no studies that combine data law and compliance in the context of redressing the imbalance or privacy between user and vendor.

In the introduction section, it was noted that Linden et al. (2020) observed that “*many [vendors’ privacy] policies still do not meet several key GDPR requirements or their improved coverage comes with reduced specificity*” [12]. This, in a way, can still be shown to be true – certainly in terms of reduced specificity. However much the vendors improve, though, there is still a fundamental problem: GDPR, and its UK counterpart the DPA, do not specify any requirement for greater transparency in how the vendors are using data for brokering advertising. Future study into the ways in which current data law such as GDPR appears to be lagging behind the rapid uptake of VAs and, in particular, the use of the data therefrom in the advertising industry could be of great benefit to the end user.

Privextractor, in its current form, is a proof of concept. Future work in the form of a comprehensive study of the way in which the vendors’ APIs, if they exist, would give Privextractor the ability to perform more comprehensive forensic extraction, for example; something that could demonstrate to the user exactly how their data is stored. A further research direction could work towards a tool that could demonstrate to the user how voice interactions with their VA can influence targeted advertising; this would be of great help in demonstrating to the user the value of their data.

VA manufacturers and vendors are likely to make changes to their privacy policies and statements as circumstances detect; as this happens, the information within Privextractor will become out-of-date and unreliable. The ideal goal, in this instance, would be for the dashboard to include a form of automatic updating – as a minimum, the tool should be aware that changes to the statements have been made. An interesting future work direction could focus on how Privextractor achieves this; real-time updates of the content within Privextractor based on the contents of the new privacy policy will be more of a challenge but would increase the utility of the tool, and the trust placed in it by users wishing to base their decisions on the information contained therein. In a similar vein, Privextractor might be made more useful with the addition of other VA vendors; we selected the four used in this study by market share but they are by no means the only VAs in use today.

Finally, as we have concluded that GDPR’s – and, by extension, the DPA’s – vagueness is failing users, we must address the ways in which it may reform. The California Consumer Privacy Act (CCPA) specifically mentions the ability for the user to opt out of having their personal data sold [56]. We have seen, however, that the specific way in which online advertising is brokered does not necessarily constitute a sale; rather, the data is exchanged on the basis that money may change hands later down the line if specific transactional parameters are met (a ‘clickthrough’). The CCPA is right, then, to address this but we may find that it has little effect on the distribution of a consumer’s data and the ability for the user to know exactly where their private information is ending up.

Data laws must cover a lot of bases in one legislation and making them too specific might be detrimental in other, unexplored ways. We feel that it is important that future work looks into ways in which the law may walk the line of being general enough to protect in all cases, but specific enough that gaps are not there to be exploited. One issue with data law is the rate at which technology advances; the current, huge rise in VA adoption took place in a few short years and any future data law must be prepared for the ‘next big thing’ that may open up whole new areas of concern for data privacy.

6 REFERENCES

- [1] Statista The 100 largest companies in the world by market capitalization. *Statista* (2022).
- [2] Amazon Amazon Advertising. *Amazon* (2022).
- [3] Apple Apple Search Ads. *Apple* (2022).
- [4] Google Google Ads. *Google* (2020).
- [5] Microsoft Microsoft Advertising. *Microsoft* (2022).
- [6] Clement, J. Advertising revenue of Google from 2001 to 2021. *Statista* (2022).
- [7] Gibbs, S. Google Nest Audio review: smart speaker gets music upgrade. *The Guardian* (2020).
- [8] V M Radhika, A. T., M Abdul Nizar An enhanced model for behavioral targeting in online advertising. *2016 International Conference on Data Science and Engineering (ICDSE)* (2016).
- [9] Hoy, M. B. Alexa, Siri, Cortana, and More: An Introduction to Voice Assistants. *Medical Reference Services Quarterly* (2018).
- [10] Edith G.Smit, G. V. N., Hilde A.M. Voorveld Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in Europe. *Computers in Human Behaviour*, 32 (2014), 15-22.
- [11] Song Liao, C. W., Long Cheng, Hongxin Hu, Huixing Deng *Measuring the Effectiveness of Privacy Policies for Voice Assistant Applications*. City, 2020.
- [12] Thomas Linden, R. K., Hamza Harkous, and Kassem Fawaz The Privacy Policy Landscape After the GDPR. *Proceedings on Privacy Enhancing Technologies* (2020), 47-64.
- [13] Kinsella, B. *UK Smart Speaker Adoption Surpasses U.S. in 2020 – New Report with 33 Charts*. City, 2021.
- [14] Tankovska, H. Number of digital voice assistants in use worldwide from 2019 to 2023 (in billions). *Statista* (2020).
- [15] Mondal, V. S. a. M. *Understanding and Improving Usability of Data Dashboards for Simplified Privacy Control of Voice Assistant Data*. City, 2022.
- [16] Zibuschka, J. A. H., Moritz AND Kubach, Michael *The ENTOURAGE Privacy and Security Reference Architecture for Internet of Things Ecosystems*. City, 2019.
- [17] Commission, E. *Privacy Flag*. City, 2022.
- [18] Aberer, H. H. a. K. F. a. R. L. a. F. S. a. K. S. a. K. *Polis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning*. City, 2018.
- [19] Erdos, D. Dead Ringers? Legal Persons and the Deceased in European Data Protection Law. *University of Cambridge Faculty of Law Research Paper No. 21/2020* (2020).
- [20] Office, I. C. s. Our history. *Information Commissioner’s Office* (2018).
- [21] Wolford, B. What is GDPR, the EU’s new data protection law? *GDPR.EU* (2020).
- [22] EUR-Lex The General Data Protection Regulation. *EUR-Lex* (2022).
- [23] ICO *The UK GDPR*. City, 2022.
- [24] Government, H. Data Protection Act 2018. *gov.uk* (2018).
- [25] Office, I. C. s. Introduction to the Age appropriate design code. *Information Commissioner’s Office* (2022).
- [26] Josephine Lau, B. Z., Florian Schaub sac Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction* (2018).
- [27] Laura Burbach, P. H., Nils Plettenberg, Johannes Nakayama, Martina Ziefle, André Calero Valdez x[S1] “Hey, Siri”, “Ok, Google”, “Alexa”. Acceptance-Relevant Factors of Virtual Voice-Assistants. *IEEE Xplore* (2019).
- [28] Frank Ebbers, J. Z., Christian Zimmermann, Oliver Hinz User preferences for privacy features in digital assistants. *Electronic Markets*, 31 (2021), 411-426.
- [29] William Seymour, M. K., Reuben Binns, Max Van Kleek *Informing the Design of Privacy-Empowering Tools for the Connected Home*. City, 2020.
- [30] Markus Lennartsson, J. K., Marcus Nohlberg Exploring the meaning of usable security – a literature review. *Information and Computer Security* (2021).

- [31] Chen Yan, X. J., Kai Wang, Qinlong Jiang, Zizhi Jin, Wenyuan Xu A Survey on Voice Assistant Security: Attacks and Countermeasures. *ACM Computing Surveys* (2022).
- [32] Emily McReynolds, S. H., Timothy Lau, Aditya Saraf, Maya Cakmak, Franziska Roesner Toys that Listen: A Study of Parents, Children, and Internet-Connected Toys. *CHI '17: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017).
- [33] Lauren N. Girouard-Hallam, H. M. S., Judith H. Danovitch Children's mental, social, and moral attributions toward a familiar digital voice assistant. *Human Behaviour and Emerging Technologies* (2021).
- [34] Yousra Javed, S. S., Akshay Jadoun x[S13] Alexa's Voice Recording Behavior: A Survey of User Understanding and Awareness. *ARES '19: Proceedings of the 14th International Conference on Availability, Reliability and Security* (2019).
- [35] Marcia Ford, W. P. Alexa, are you listening to me? An analysis of Alexa voice service network traffic. *Personal and Ubiquitous Computing*, 23 (2019).
- [36] Alex Akinbi, T. B. Forensic Investigation of Google Assistant. *SN Computer Science* (2020).
- [37] Hyunji Chung, J. P., Sangjin Lee Digital forensic approaches for Amazon Alexa ecosystem. *Elsevier*, 22 (2017), S15-S25.
- [38] Bhupendra Singh, U. S. A forensic insight into Windows 10 Cortana search. *Computers & Security* (2017), 142-154.
- [39] Horsman, G. Loose-Lipped Mobile Device Intelligent Personal Assistants: A Discussion of Information Gleaned from Siri on Locked iOS Devices. *Journal of Forensic Science*, 64 (2019).
- [40] Lareo, X. TechDispatch #1: Smart Speakers and Virtual Assistants. *European Data Protection Supervisor* (2019).
- [41] Irwin, L. What is an ISO 27001 risk assessment and how should you document the process? *itgovernance.eu* (2020).
- [42] Wood, S. Blog: Using biometric data in a fair, transparent and accountable manner. *Information Commissioner's Office* (2019).
- [43] Kemp, R. Big data and data protection (GDPR and DPA 2018). *Reuters Practical Law* (2020).
- [44] Samonte, M. Google v CNIL Case C-507/17: The Territorial Scope of the Right to be Forgotten Under EU Law. *European Law Blog* (2019).
- [45] BBC Amazon Alexa security bug allowed access to voice history. *BBC News* (2020).
- [46] Cook, J. Amazon employees listen in to thousands of customer Alexa recordings. *The Daily Telegraph* (2019).
- [47] GCHQ Investigatory Powers Act. *GCHQ* (2019).
- [48] Economics, L. S. o. Could the European GDPR undermine the UK Investigatory Powers Act? *London School of Economics* (2016).
- [49] Daniel J. Dubois, R. K., Anna Maria Mandalari, Muhammad Talha Paracha, David Choffnes, Hamed Haddadi When speakers are all ears - Understanding when smart speakers mistakenly record conversations. *20th Privacy Enhancing Technologies Symposium (PETS2020)* (2020).
- [50] Amazon Set Up Alexa Hands-Free on Your Phone. *Amazon* (2020).
- [51] Waranch, R. S. Digital Rights Ireland Deja Vu: Why the Bulk Acquisition Warrant Provisions of the Investigatory Powers Act 2016 Are Incompatible with the Charter of Fundamental Rights of the European Union. *George Washington International Law Review* (2018).
- [52] Cranor, P. E.-N. A. H. D. A. Y. A. A. L. F. *Exploring How Privacy and Security Factor into IoT Device Purchase Behavior*. City, 2019.
- [53] eMarketer Amazon's ad revenue in 2020 is set to grow 23.5% despite the pandemic. *Business Insider* (2020).
- [54] Burgess, M. This is how age verification will work under the UK's porn law. *Wired* (2019).
- [55] Jenny Radesky, Y. L. R. C., Nusheen Ameenuddin, Dipesh Navsaria Digital Advertising to Children. *American Academy of Pediatrics* (2020).
- [56] *California Consumer Privacy Act (CCPA)*. State of California Department of Justice, City, 2022.