

University of Oxford
Oxford OX1 2JD | United Kingdom



UNIVERSITY OF
OXFORD

Here's Charlie!

The Semantic Web Vision of Agents in the Age of LLMs - and steps we need to take in semi-automating data governance to get there

By:
Jesse Wright

1. Introduction

This paper presents our research towards a near-term future in which *legal entities*, such as *individuals* and *organisations* can entrust semi-autonomous AI-driven agents to carry out online interactions on their behalf. The author’s research concerns the development of semi-autonomous Web agents, which consult users if and only if the system does not have sufficient context or confidence to proceed working autonomously. This creates a user-agent dialogue that allows the user to teach the agent about the information sources they trust, their data-sharing preferences, and their decision-making preferences. Ultimately, this enables the user to maximise control over their data and decisions while retaining the convenience of using agents, including those driven by LLMs.

We discuss why introducing such agents is crucial to support user autonomy on the Web and in society; and explore how the advent of LLMs enables the advancement of general-purpose Web agents. In view of developing near-term semi-autonomous agents, we address the question: “How do we build a trustworthy and reliable network of semi-autonomous agents which represent individuals and organisations on the Web?” In particular, presenting key requirements for supporting safety guarantees around *belief*, *data sharing* and *data usage* whilst enabling *serendipitous* dialogues between software agents.

In the second half of this paper, we focus on the challenges of semi-automated *data sharing* and present a sociotechnical vision working towards semi-automated data sharing between semi-autonomous Web agents.

2. Background: why agents, and why now?

2.1. Should we use agents?

As Wooldridge [1] discusses, it is common to “oversell agent solutions, or fail to understand where agents may usefully be applied” and thus it is important to have “realistic expectations of what agent technology can provide.” Consequently, it is reasonable to expect that current excitement around many LLM-based agents is overhyped¹. We are working towards the introduction of semi-autonomous agents that faithfully represent different legal entities on the Web. Thus, each agent represents different interests and thus have differing *objectives* and *controls* as well as access to differing corpus’ of private data. This aligns well with Wooldridges proposed use of agents when there is a “Distribution of data, control or expertise.”

2.2. The evolution of agentic capabilities and the impact of LLMs

2.2.1. Defining agents and understanding agentic communication

Abstractly, an agent is “a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its design objective” [1, p. 15; 2]. Just as humans must negotiate and cooperate to achieve shared goals, so too must agents within multi-agent systems [1, p. 24-25]. This is only possible if agents have an effective means

¹<https://www.kadoa.com/blog/ai-agents-hype-vs-reality>

of communication. To communicate they must have a **(R1)** shared conceptual understanding of the topic on which they communicate (e.g. two agents can only communicate about weather if they both ‘know’ what it means to be sunny, raining and overcast, and have a notion of temperature and a measure - such as degree celsius - by which to describe it) and **(R2)** a means of sender (resp. receiver) encoding (resp. decoding) messages to transmit these concepts between agents [3]. We shall refer to the requirement that agents are “capable of autonomous action [execution]” [1, p. 15] as **(R1E)**.

2.2.2. Information systems and APIs

Since the emergence of autonomous agents in the 1980’s [1, p. 304] there have been numerous developments in the way in which agents understand **(R1)** and transmit **(R2)** concepts.

First, came systems where human programmers would encode specific execution semantics in code **(R1E)** for instance the “concept of temperature” would be defined by a function which takes a reading from a temperature sensor and these systems would communicate **(R2)** using a structured encoding such as sending JSON encoded objects over the Web - containing a “temperature” key. In this paradigm, the conceptual understanding (e.g. of temperature) **(R1)** did not fully live within the system, but was rather implicitly communicated between systems in the form of API documentation where the developers implementing each system would explain to each other “the temperature key in the JSON object is a 32bit floating point value representing the temperature, in central London, as measured in degrees celsius”. Naturally, these primitive agents had very limited *versatility* as human developers were required to define the execution semantics and conceptual semantics (documentation) by one exact concept at a time. Moreover, these agents, which are still widespread in the form of Web APIs, face an *interoperability* problem - as there are countless Web APIs all of which have a “temperature” key in the JSON object they return; but with different meanings due to the implicit semantics of the metric, collection location, collection date etc. encoded in the API documentation. This tightly couples the client-agent to a single server agent and to communicate with more “agents” the developer must manually write code to decode the different encodings of different server agents and align the conceptual understanding across different documentation.

2.2.3. Semantic Web agents

Subsequent developments saw attempts to migrate this conceptual understanding **(R1)** from documentation to data which is sent between agents [4, 5, 6, 7]. In domains such as the Semantic Web, conceptual understanding is described symbolically using ontologies [1, p. 180] and rules [8, 9] **(R1)** - and encoded in highly generic RDF [10] syntaxes **(R2)**. With these agents (systems) [3] possessing a symbolic understanding of concepts, they are able to precisely describe and reason about the information they send and receive. This increases agent *interoperability* and *portability* by making explicit the semantics that are often implicitly encoded in API documentation; moreover, agent *versatility* increases with the ability to describe requests and responses on the fly without being constrained to use the finite set of concepts listed in API documentation. However, ontologists are still required to build the vocabularies and rules for concept description (e.g. someone must manually define “degrees celsius” as a “unit of measure”

for “temperature”; someone must provide the rules that define the mapping between “degrees celsius” and “fahrenheit” so that systems using two different measures can interoperate) (**R1**) and developers are still required to write code that defines the execution semantics (**R1E**) that is, a system can semantically formulate the question “what is the current temperature in London in degrees Celsius”; but a developer is then still required to write the code that then fetches the temperature from the sensor. Furthermore, in this paradigm “If two agents are to communicate about some domain, then it is necessary for them to agree on the terminology that they use to describe this domain.” [1], often requiring agents to share the same vocabularies and hence “worldviews”.

2.2.4. LLM-Powered Agents

Now LLM-powered agents are emerging, with many using LLMs trained upon a textual corpus containing a vast array of human knowledge. Consequently, these agents have been demonstrating a greater *breadth* and *depth* of conceptual understanding (**R1**) than human developers could imagine encoding by hand in formal ontologies. Moreover, LLMs can encode and decode these concepts in natural language (as well as, increasingly, machine syntaxes [11]) (**R2**). Conversational LLMs also possess inherent execution semantics (**R1E**) with the capability to formulate a written response to any input which they are given. These execution semantics, however, do not extend to access of system-level knowledge or resources (such as temperature sensors) and are less reliable/deterministic than production-ready implementations of previous generations of agents. The trade-off of the breadth and depth of LLM understanding; is an increased vagueness of concepts and the presence of internal inconsistencies in LLM conceptualisations - moreover, there is no single worldview, per se, that the language model holds; and, just as with humans, the facts that an LLM claims are highly dependent on the context of the conversation in which they occur.

2.2.5. Neurosymbolic Agents

Hybrid LLM and Semantic Web technologies have emerged [12, 13] as a branch of Neurosymbolic AI [14]. These technologies leverage the strengths of Large Language Models (LLMs) discussed in Section 2.2.4 along with their intuitive, System 1 reasoning [15]. These hybrid technologies also benefit from the precise, logical System 2 reasoning [15], safety guarantees, and structured knowledge representation offered by Semantic Web technologies. Neurosymbolic AI-backed agents are becoming increasingly prevalent [16], and we anticipate that the protocol outlined in Section 3 will primarily be implemented by such agents.

2.3. Agents for supporting user-autonomy on the Web

2.3.1. What are the problems?

We posit that semi-autonomous Web agents have the potential to help us escape the perversions of the modern Web² [17] and create a world in which technology un-obtrusively facilitates users’ lives in the manner in which they desire. On the Web today, numerous issues around

²<https://www.cnn.com/2019/03/11/tim-berners-lee-the-web-is-dysfunctional-with-perverse-incentives.html>

user experience and user autonomy arise from the fact that the services that users interact with are not aligned with the *interests* or personalised to the *preferences* of users. Yet these service providers are responsible for the content and presentation of what users interact with - through the Websites and applications that they build. In his letter “30 years on, what’s next #ForTheWeb?” Berners-Lee identified 3 sources of dysfunction on today’s Web

30 years on what’s next #ForTheWeb?

1. **Deliberate, malicious intent**, such as state-sponsored hacking and attacks, criminal behaviour, and online harassment.
2. **System design that creates perverse incentives** where user value is sacrificed, such as ad-based revenue models that commercially reward clickbait and the viral spread of misinformation.
3. **Unintended negative consequences of benevolent design**, such as the outraged and polarised tone and quality of online discourse.

Additionally, the rising cognitive load associated with information discovery [18] and the proliferation of busy work—such as bureaucratic administration—driven by an always-online culture, have been shown to be detrimental to society. As Vuori et al. note, “As advanced and ‘intuitive’ information systems ‘enable’ individuals to operate them for trivial tasks, experts now perform tasks that have traditionally been the core of assistants’ work (e.g., reporting, travel booking, etc.)” Previous works have already applied personal agents in a range of scenarios, such as virtual assistants aimed at reducing cognitive load and enhancing decision-making [20, 21, 22].

2.3.2. What remedies to agents offer?

Although many of the aforementioned challenges have deep social roots—such as greed and economic incentives [23]—we contend that Web agents, when genuinely acting in the best interests of users, have the potential to reshape how these social and economic dynamics manifest. This belief is grounded in several assumed *values* and *capabilities* of such agents, which we recognize may not be fully realized due to technical and social constraints. Our assumptions are as follows:

1. **User Representation:** Agents will accurately and faithfully represent users’ interests and intentions in the online environment,
2. **Trust and Preferences:** Agents will adhere to user-defined preferences in determining which signals to use when evaluating whether other agents or services should be trusted for specific purposes—whether as authorities on certain topics, as custodians of user data, or as service providers. These signals may include social relationships, trust in particular institutions (e.g., government entities), endorsements from other agents, past interactions, and legal agreements or regulations that stipulate penalties for misconduct.
3. **Autonomy and Delegation:** The degree of autonomy and control delegated to the agent will be user-defined and adjustable at any time. This could range from complete delegation, such as “organize all my meals, meetings, travel (including holidays), and exercise plans using all relevant data from my knowledge base and IoT devices without

prior consultation,” to specific instructions like “book my flights for next week, considering my stated sustainability and budget requirements, and suggest some gyms and workout routines tailored to my longevity goals based on recent health research,” to precise commands such as “book flight HC123 on Oct 12 for me.” This balance allows users to reclaim cognitive resources while maintaining an appropriate level of oversight, depending on their individual preferences.

4. **Interoperability:** Agents will be interoperable with a wide range of Web services and with other users’ and organizations’ agents, ensuring equitable access to online services through these agents.
5. **Fair Representations:** Agents will be capable of advocating for user interests in an open marketplace, ensuring fair competition and representation.
6. **Accessible Interaction:** Agents will be *accessible* to a diverse range of user groups and available through various interaction modalities—such as mobile apps, voice assistants, smart-home devices, or landline calls—adapted to the user’s needs and capabilities.

How can agents reshape the influence of various socioeconomic forces on the Web? Consider the current monetization model that dominates the digital landscape: advertising. Many social media platforms have discovered that the most effective way to maximize profits under this model is to capture as much of the user’s attention as possible—thereby increasing the number of ads they can display and influence users with. While the predictable consequence is the loss of users’ time, the less predictable, but arguably more harmful, effects include increased polarization [24] and mental health issues [25] among the user population.

Web agents have the potential to dismantle these attention [26] and surveillance [27] economies and move us towards an intention economy [28], where agents negotiate for products and services on our behalf, based on the preferences and policies we set as users. Several factors can drive this transformation, most importantly the role of agents as intermediaries between users and services on the Web.

Instead of users having to accept the provider’s choice over content and its presentation, semi-autonomous agents can filter and prioritize information that aligns with the user’s intentions and needs, delivering only what is relevant and valuable.

At the same time, these agents have access to enough user data to accurately identify products that match individual needs and provide recommendations based on a genuine assessment of their quality. This offers a remedy to the current bias of consumers to select products and services that are best *marketed* or *influenced* rather than those that best match their needs [29, 30].

This also allows users to bypass many of the algorithmic harms [31] that are pervasive in today’s platforms, fundamentally changing the dynamics of how digital services interact with their audience.

3. A protocol for Agentic Communication at Web Scale

As discussed in Section 2.2 there exists a substantial body of research on communication between multi-agent systems. The vision of the Semantic Web itself [32, 33, 34] as shown by Charlie, the “AI that works for you” was centred around Web Scale agentic systems. Yet, the 2006 lamentation that “[b]ecause we haven’t yet delivered large-scale, agent-based mediation, some commentators argue that the Semantic Web has failed” [35] still rings true today. The growing use of LLMs raises a key challenge in building Trustworthy and Reliable Web Agents [36, 37]. In this section we present a communication protocol between Neurosymbolic AI Agents at Web Scale.

3.1. Design Requirements

We identify the following non-functional requirements for an agent communication protocol. It must be possible for semi-autonomous agents to:

1. *Identify* legal entities, such as individuals or organisations, on the Web [38] so they can be referenced.
2. *Deterministically discover* other agents representing an entity from their Web identity [38]. This does *not* require all agents to be publicly advertised; some may be discovered from links to protected documents.
3. Describe, and agree to, any *usage controls* [39, 40, 41] associated with data they exchange. This allows sharing of protected data while articulating the recipient’s legal or moral obligations [42].
4. Describe the *origin* and *provenance* of data they exchange. In an open world of agents that can “say anything about anything,” systems can identify which external claims to believe for a given task, based on the agent’s internal trust model.
5. *Unambiguously* describe *ground truths* they send, and *agreements* they make, using a formal representation. Consider the case where an individual’s agent purchases a flight from an airline’s agent. Structured ground truths eliminate an LLM’s risk of hallucination or misinterpretation of key information, such as the flight time (“10 o’clock” could be 22:00 or 10:00). As agents represent entities in binding agreements, this approach also reduces the risk of legal disputes by limiting the subjectivity of agreed terms and thus the ability to reinterpret or rescind them [43]. Furthermore, agents can implement rule-based internal safeguards, such as user-defined daily spending limits. Truly generic agents may generate and communicate structured ontologies when encountering new tasks. In many cases we expect LLM-supported ontology construction [44] to facilitate generation; however, research is required to understand how (1) agents can align on conceptual models for use and (2) how human oversight can be maintained without disrupting user experience.
6. Contextualise a task which may be *ambiguous* or poorly defined, such that interacting agents can introduce new solution spaces or negotiating actors in a *serendipitous* manner.

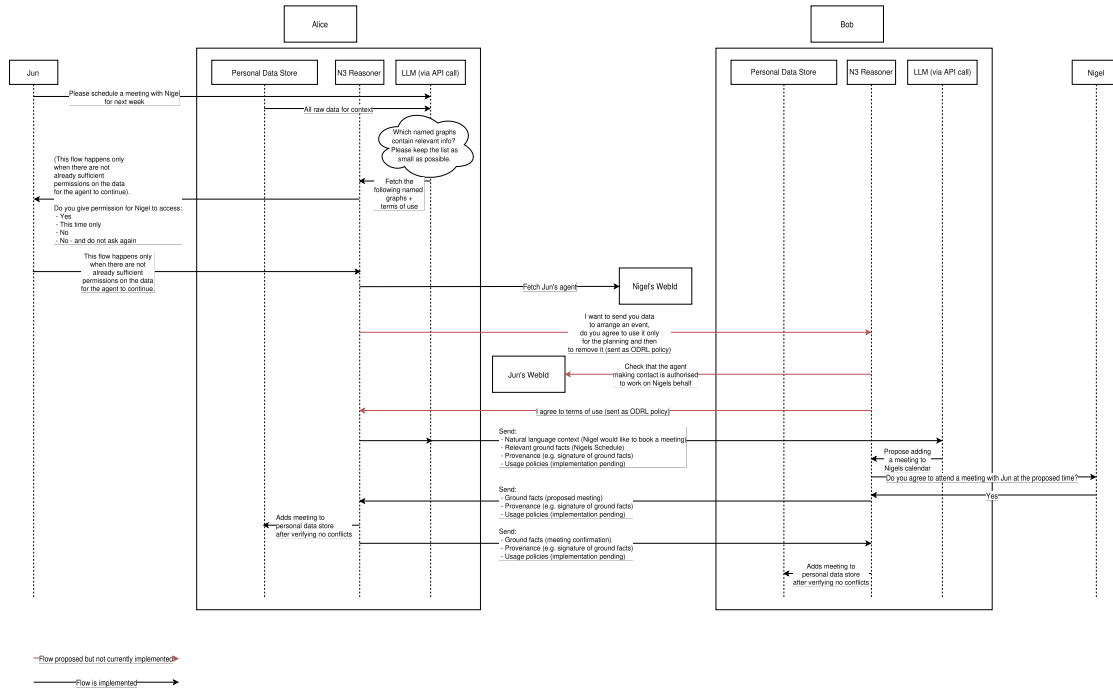


Figure 1: Flow diagram for the scheduling use case. Alice is Jun’s agent and Bob is Nigel’s agent.

3.2. Sample Use-Case and Implementation

We implemented the following flow where agents act as personal assistants for individual users:

1. Jun types into a chat “Schedule a meeting with Nigel next week”;
2. Jun’s agent identifies data to be shared with Nigel and requests relevant sharing permissions from Jun (where not already obtained);
3. Nigel’s agent receives a request from Jun;
4. Nigel is prompted to confirm that he believes Jun is an authoritative source of truth for her calendar (where not already obtained);
5. Nigels agent proposes a meeting time to Nigel; and
6. the meeting is proposed to Jun’s agent and automatically confirmed.



We have created a running demo with a video, flow-diagrams (including Figure 1) and other resources for our codebase³. The implementation corresponds to the above use-case steps:

1. Given the user prompt and a set of known WebID profiles [38], an LLM called by Jun's agent identifies the relevant entities for the agent to negotiate with (Nigel), and the WebIDs of those entities. Given the user prompt, and the user's personal knowledge graph, an LLM called by Jun's agent identifies which subset (as a list of named graphs) of the user data are needed to fulfil the user's request.
2. Notation3 [45] reasoning is used to identify the policies applicable to the data subset. In the available demo recording, policies are encoded in ACP [39]; we are currently migrating to use ODRL [40] and DPV [41]. If these policies do not yet permit read access to Nigel, Jun is prompted to modify them. Jun's agent then dereferences Nigel's WebID [38] to discover information about his agent.
3. Jun's agent uses an LLM to construct a message for Nigel's agent, explaining the context of Jun's task: "Jun seeks to schedule a meeting for next week. Propose a time for Jun and Nigel to meet using their calendars." Jun's agent sends Nigel's agent this message along with the RDF description of Jun's calendar and any associated policies and provenance. With ACL, Nigel's agent does not need to agree to any policy obligations; this changes with ODRL. The provenance in this case is simply a signature of the canonicalised calendar dataset [46] using Jun's public key.
4. As Nigel has instructed his agent that Jun is an authoritative source of information on all topics, his agent *believes* (takes as ground truth) the signed RDF dataset sent by her agent. We are developing conceptual models for agentic trust; these extend existing trust vocabularies [47, 48, 49, 50] with a range of features including (1) qualifying whether sources are trusted for particular *types* of claims; for instance, most agents should trust certified airlines to present flight times and prices, but not medical data (2) qualifying the forms of provenance *secure* enough for a given task; for instance, an insurance provider may require provenance demonstrating a user was signed in with two-factor authentication when entering financial details to their knowledge base.
5. Nigel's agent proposes a meeting time, using the natural language context (*not* a ground truth) and the calendar dataset (ground truth). The LLM proposes a meeting time, then the N3 reasoner applies rules to (1) ensure no calendar conflicts and (2) check for user confirmation, before adding the proposed time to the knowledge base. In a future iteration, we plan to use the LLM to generate an N3 query that proposes a meeting time based on Nigel's Personal Data Store and Jun's calendar.
6. Upon meeting the above requirements, the reasoner sends to Jun's agent a meeting proposal, in the form of an RDF dataset with attached usage policies and provenance. Jun's agent confirms this dataset can be believed based on the internal trust model. The rules within Jun's agent validate that there are no conflicting events. Jun's personal knowledge base is updated with the event, and a confirmation is sent to Nigel's agent.

³<https://github.com/jeswr/phd-language-dialogue-experiment>

4. Semi-automated data sharing for the Web, Web Agents and Data Spaces

In the ever-evolving landscape of digital privacy, the management of personal data, including cookies, within Web browsers has become increasingly crucial. Legislative attempts to give users back control over data that is captured in cookies have largely resulted in obstructive cookie notice pop-ups across the Web containing convoluted or misleading policies, which new research suggests are not often compliant with user preferences [51]. Consequently, cookie policies and other Terms of Service agreements are deemed “biggest lie on the internet” [52].

In this section, we present our vision of how the Open Digital Rights Language (ODRL) [40, 53], Data Terms of Use (DToU) [54] and the Data Privacy Vocabulary (DPV) [55] can be embedded into websites with RDFa [56], and transmitted within HTTP Headers in order to well-describe and allow negotiation over the terms of use that are applied to cookies in the browser. We argue that if deployed alongside regulatory incentives, or pressure, this technology stands to benefit individuals, industry and regulators. **Individuals** stand to benefit from a smoother experience and enhanced, semi-automated control over their privacy on the Web, with browsers managing cookie policies on their behalf. **Businesses** stand to gain significantly from this standardised approach to cookie management. By adhering to machine-interpretable standards endorsed by regulators, companies can reduce the risk of privacy-related lawsuits, demonstrate compliance with regulations such as the General Data Protection Regulation (GDPR) [57], the ePrivacy Directive [58] or the California Consumer Privacy Act (CCPA) [59], and streamline the implementation of privacy policies. Furthermore, it becomes easier for these bodies to implement automated auditing systems that validate their compliance with their advertised terms of use. **Regulators** stand to benefit from a unified framework for describing regulations around personal data in cookies, and automated techniques for checking compliance with that regulation.

We view this work as a critical stepping stone to achieve semi-automated data governance in emerging data-centric technologies that form the next generation of the Web, such as Web Agents [60] and Data Spaces like Solid [61, 62, 63]. In particular, browser cookies are a mature, widely used and well-understood Web technology, and the measures that websites must take to handle cookies whilst complying with jurisdictional data protection regulations are well-known. This makes browser cookies a good target use case where academia, regulators and industry can ‘battle-test’ and mature technologies such as ODRL, DToU and DPV to support real-world requirements for semi-automated data governance. Thus we propose cookie data as the starting point for the rollout of annotated terms of use at Web scale.

4.1. Related Work

4.1.1. Privacy Policies in Browsers

The concept of privacy policies in browsers is not new. The W3C’s Platform for Privacy Preferences (P3P) [64] was a protocol published in 2002. P3P enabled Websites to describe their data management practices to users in a standardised and machine-readable format. This included a description of the type of data that was collected via different browser interactions,

and the purpose for which the Website was collecting it. P3P enables Websites to encode their privacy policies in XML, which can then be automatically retrieved and interpreted by Web browsers and other user agents. This allows users to easily understand a Website's data collection and usage practices without having to read convoluted privacy policies. Moreover, it enabled the browser to block parts of the Website until users had opted into certain types of their data being collected. Despite its innovative approach, P3P faced challenges in adoption and implementation, leading to limited use and eventual obsolescence as privacy concerns and regulations evolved [65].

Global Privacy Control (GPC) [66, 67] and its predecessor Do Not Track (DNT) [68] take an all-or-nothing approach to improving user privacy on the Web, introducing a binary signal which users can enable to indicate they do not want their data to be sold or shared. Unlike DNT, GPC has gained traction because it is backed by the CCPA [59, 69]. This regulation requires companies to honour user preferences of opting out of data sharing, giving GPC more legal weight [70].

4.1.2. Vocabularies for expressing cookie preferences

Bushati et al. [71] introduced the *OntoCookie* ontology, developed as part of a study investigating users' awareness of the data they consent to sharing via cookies. *OntoCookie* is a formal representation of the cookie domain, comprising 229 axioms, 32 classes, 10 object properties, and 10 data properties. It models various types of cookies, their metadata, and their purposes (e.g., necessary, analytics, marketing) using a top-down ontology engineering approach. By leveraging this ontology, Bushati et al. [71] created a KG-based tool to enhance user comprehension of cookie data sharing, providing a more transparent and interpretable view of cookie data.

Of the 32 classes introduced by *OntoCookie*, 8 are subclasses of purpose for processing cookie data (*Analytics*, *Marketing*, *Profiling*, *ServiceOptimisation*, *ServicePersonalisation*, *ServiceProvision*, *Tracking*), 10 are related to types of cookies (*Authentication*, *HostOnly*, *HttpOnly*, *Persistent*, *SameSite*, *Secure*, *Session*, *Super*, *Tracking*, *Zombie*) and 2 are subclasses of necessity, i.e., *Necessary* and *Optional*. Concerning the purpose classes introduced, only *Tracking* does not have an equivalent concept in DPV 2.0 (which otherwise has 88 additional purpose classes compared to *OntoCookie*), which may be worth adding as an extension to DPV under `dpv:Marketing`. The cookie classes are useful for adding additional descriptions about cookies, but do not help describe their terms of use, and are thus not as useful to this work.

Thus, we propose the use of ODRL, DToU and DPV in favour of *OntoCookie* as these vocabularies are better suited for describing terms of use in this use case, and better generalise to terms of use descriptions outside of the context of cookie policies which is the long-term end goal of this work.

4.1.3. Deceptive patterns and extensions for managing cookies

Deceptive patterns in cookie consent popups manipulate users into agreeing to data collection. These tactics often violate GDPR principles requiring informed and voluntary consent. Common deceptive patterns include: (1) Pre-selected Options – banners with pre-ticked boxes for non-

essential cookies [72], contrary to GDPR guidelines requiring active consent [73]; (2) Deceptive Button Colours – highlighting the “Accept” button more prominently than the “Reject” button to influence user choice [72]; (3) Complex Navigation – making users navigate multiple layers to be able to reject cookies, while accepting them is straightforward [74]; (4) Misleading Labels – declaring marketing cookies as essential to imply users cannot opt out without affecting functionality [74]; (5) Hindering Withdrawal – making it difficult to withdraw consent by not providing easily accessible options [75]; and (6) Manipulative Language – using vague or biased language to emphasise benefits of accepting cookies while downplaying data collection [76].

A study by Bollinger et al. [77] identified widespread GDPR violations across nearly 30,000 Websites, with 94.7% of these sites exhibiting at least one potential violation. This underscores the need for a technical infrastructure, developed in collaboration with regulators, that facilitates platforms in developing legally compliant terms of use.

To combat such problems, there have been several efforts to implement extensions which automate the process of accepting or rejecting cookies in browsers. For instance CookieBlock [77] is a browser extension designed to automate cookie consent management by using machine learning to classify cookies based on one of four purposes: Strictly Necessary, Functionality, Analytics, and Advertising/Tracking. The extension then automatically accepts or rejects cookies based on which of the four categories users enable. For classification, CookieBlock achieves a mean validation accuracy of 84.4% and filters out approximately 90% of privacy-invasive cookies without significantly affecting Website functionality [77]. This approach does not depend on the cooperation of Websites, thus improving user privacy even on sites that do not comply with GDPR requirements.

As Bollinger et al. [77] acknowledge, CookieBlock, while innovative, is a temporary client-side fix. No solution limited to the client can address the root problems of (1) Websites presenting ambiguous and convoluted privacy policies that may be misinterpreted by humans and machines, and (2) Websites ignoring user consent when personal data is allowed to be used for a limited set of purposes. By nature of the CookieBlock being ‘adversarial’, it broke 10% of websites. Our proposal in Section 4.2 avoids this problem, with Websites able to describe cookie configurations required for Website functionality. Our proposal also offers more fine-grained user preferences, offering a wider range of purposes and controls for properties including the cookie retention period.

4.2. Vision for semi-automated data sharing on the Web

In our long term vision of the future Web, all data shared between clients (including browsers and other types of agents like those discussed in Section 3) and servers is annotated with machine-readable terms of use agreements. This enables the data sender (server or client) to declare features such as which legal basis is being used to process said data; and data subjects to express what permissions, obligations and other restrictions apply to the usage of their data. In turn, this enables the data recipient (client or server) to automatically and unambiguously determine how the data may or may not be used, using rules-based reasoning. The data recipient may also use terms of use *requests* to identify their promises and the permissions they would like the sender to include in the agreement. We expect these machine-readable terms of use

agreements to be encoded using RDF [10], described using the ODRL [40] or DToU [54] models and to extend upon terms from DPV [55].

We now turn specifically to the case of cookie management, which is the focus of this paper. We envision migrating towards a state where the act of a data subject (in this case the user of a browser) ‘permissioning’ to the processing and sharing of personal data within cookies is communicated by having browsers including accompanying ‘Data-Policy’ header(s), containing terms of use agreements, for each HTTP Cookie header (c.f. RFC 6265) in the browser request. The contents of the terms of use agreements in the ‘Data-Policy’ header(s) are to be generated by the browser or a browser extension enforcing the users’ data sharing preferences against the Website’s request. As we shall elaborate further in the remainder of this section; this process of enforcement may either be performed statically in the browser or involve a ‘negotiation’ between the Website and browser agent. This process may require explicit user input depending on what existing privacy preferences the user has supplied and the legal basis used by the Website for the processing of cookie data, e.g., requiring explicit GDPR consent from users will imply their affirmative and freely given acceptance of the cookies’ terms of use.

In the following sections we discuss a 3-step pathway towards reaching the vision, where each step itself is a valid technical solution, gradually increasing the automation and formality.

4.2.1. Step 1: Machine readable terms of use requests in cookie dialogues

As a first step towards this goal, we propose that Websites begin embedding terms of use requests as machine-readable RDFa [78] in existing cookie dialogues. In particular, for each cookie identifying the type of data that is retained and a granular description of the purposes for which it will be processed. Further, we propose the use of a standardised naming scheme for HTML `id` attributes for the checkboxes to opt in or out of particular cookies. This would enable browser extensions to automatically manage the selection boxes on the behalf of users in a deterministic manner. This automated selection would be made based on clear preferences that the user has actively chosen. Unlike existing browser extensions which accept/reject cookies [77], this solution does not rely upon Large Language Models (LLMs) or other heuristic measures to interpret the natural language cookie policies displayed, and accept/reject cookies based on broad categories of cookies such as performance, functional and analytics. Instead, this solution allows for matching against well-defined and fine-grained user preferences in the browser.

Our proposal using RDFa embeddings would be relatively straightforward for existing Consent Management Platforms (CMPs) to independently roll-out as they are already in control of the current pop-up functionality. Having a browser extension interact with the existing HTML element, rather than posting data directly to an API also means that there is no interim effort required to align the consent management APIs that CMPs use.

Similarly to Bollinger et al. [77], limited client (browser) side enforcement of user preferences may be introduced once embedded RDFa descriptions are available. In particular, the browser may prevent the creation of cookies which may not be used for any purposes according to user preference. At this stage of the work this simply means that the only cookies permitted on a given Website are those that have been ‘allowed’ by the extension.

For those websites that do not move to use RDFa embeddings to describe the purposes;

there is also the option of using LLMs called by a browser extension in the short-term to translate the natural language privacy policies into machine-readable terms of use requests. These formal descriptions can then be used by browser extensions to automatically accept or reject cookies based on user preferences. This LLM-based supplement for RDFa embeddings is an ad hoc solution in that it does not address the root problem of ambiguous and complex natural language privacy policies, but rather provides a temporary fix to the problem of manual cookie management. However, it would improve the granularity of control that users have in comparison to existing solutions such as CookieBlock [77].

4.2.2. Step 2: Machine-readable terms of use requests in headers

Over time, we recommend that Websites also include these terms of use requests as HTTP headers with the name ‘Data-Policy-Request’, which contains either the full privacy policy for Website cookies, links to a cacheable description of the privacy policy, or links to an API for policy negotiation (c.f. Section 4.2.3). Whenever both a browser and Website implement the ‘Data-Policy’ detailed in Section 4.2.3, there would no longer be a need for server-managed cookie user interfaces.

Unlike Section 4.2.1, even Websites using CMPs will need to take individual action to implement this header-based proposal, as CMPs do not generally intercept response headers. Consequently, adoption is likely to be slower.

4.2.3. Step 3: Consent and machine-readable terms of use agreements

In parallel to the recommendations of Sections 4.2.1 and 4.2.2, we suggest browsers start implementing the ‘Data-Policy’ header. This allows browsers to communicate to Websites, on behalf of users, the permissions that users have given for the processing and sharing of their personal data collected in cookies. This role is currently fulfilled by Consent Management Platforms (CMPs), which use custom APIs and flows to collect and log user consent.

This proposal is backwards compatible with HTTP servers and browser clients that do not recognise ‘Data-Policy’ headers as “a proxy MUST forward unrecognised header fields” and “other recipients SHOULD ignore unrecognised header and trailer fields.” [79]. We propose the name ‘Data-Policy’ rather than ‘Cookie-Policy’ in anticipation that the header will also transmit terms of use agreements for other non-cookie headers and the message body in the future.

With ‘Data-Policy’ headers, browser clients can enforce terms of use agreements which do not permit the cookie data to be used for any purpose, by blocking the cookie from being sent. Clients will not be able to enforce terms of use agreements that allow the Website owner or third parties to use personal data for a limited range of purposes. In these cases the Website and third parties are responsible for adhering to the agreed terms of use; similarly to how Websites and third parties are expected to respect consent signals sent out by CMPs today. Rather than relying on platforms benevolently respecting these terms of use agreements, our goal is to work with regulatory authorities to make these terms of use agreements legally enforceable. In EU countries and the UK, we can start by making ‘Data-Policy’ headers the recommended way to signal legal consent to process personal data. In the future, we envision these terms of use agreements being of a more contractual nature, similar to a data sharing agreement. As we

shall re-iterate in the following subsections and Section 4.5, this is where a joint approach is required between research, regulatory bodies and industry, in order to concurrently develop technical standards, new regulations and best practices to:

1. Ensure that with this proposal, companies can comply with EU and UK regulation for consent collection and logging [80] which is currently fulfilled by custom flows that take place when users confirm their preferences in cookie consent dialogues.
2. Provide regulatory incentives or pressures for adoption of the ‘Data-Policy’ header. Incentives could include “Safe Harbour”⁴ [81] clauses in data protection regulation, which legally protect companies using policy evaluation engines to ensure terms of use agreements are respected. As a more stringent measure, regulations could mandate the use of such policy engines and require companies to undergo regular audits to verify their compliance and obtain certifications from Data Protection Authorities.
3. Align this proposal with existing semi-automation approaches adopted for data governance within enterprises [82].
4. Ensure this proposal enables the reduction of long-term engineering, legal and other compliance-related costs that enterprises face [83].

4.2.4. Negotiating terms of use agreements

We anticipate use cases where terms of use agreements cannot be immediately computed by browsers using static terms of use requests from Websites and user preferences (that can be either pre-defined or obtained from the user in real-time via a browser-managed popup). For instance, some Websites offer a choice between paying for services and consenting to the use of personal data for targeted advertising [84]; others may present ads at a lower frequency when targeted advertising is possible [85]. Such Websites may need to offer a ‘negotiation’ API where browsers can perform operations such as payments and obtain a service contract, complementing their terms of use agreement, that allows use of the Website without sharing data for targeted advertising.

As discussed by Solove [86] and Florea and Esteves [87] obtaining valid GDPR consent remains an issue, as it implies the user knows the purpose for which their data is being used, as well as the identity of the legal entity ‘behind’ the processing of said data, among other conditions. As such, the development of a policy-based Web environment must not shy away from going beyond consent to explore other legal bases, while, of course, relying on it as an information safeguard for users.

4.2.5. Beyond Cookies

We are striving for a future in which all data sent over the Web is annotated with terms of use agreements between the sender and recipient. This could be achieved by having HTTP requests and responses always containing the ‘Data-Policy’ header. This header would contain terms of use agreements that would be applicable for data in the headers and message body with the

⁴A safe harbour provision is a legal clause that offers protection from liability or penalties under specific conditions, as long as the party complies with certain predefined guidelines or standards.

possibility to have fine-grained definitions of different terms of use for different parts of the request or response. This enables the sender to be explicit about any requirements they have for how their data is governed, and for the recipient to implement policy engines that ensure these requirements are respected. In the worst case, if a server is not able to understand or handle the terms of use of an incoming HTTP request, we would expect the server to return a 5xx response and discard of any user data received in the request. It would be best practice for a server to also indicate any modifications required to the terms of use in order for a future request to be accepted, using an RDF-encoded response. If a client receives a set of terms of use that it is not able to respect, it should also immediately discard of the information received in this response.

4.2.6. Beyond Websites

In the spirit of the Semantic Web [60, 33, 34], we posit that over time the Web will evolve away from users sending and receiving data via Websites, to having their interactions on the Web mediated via Web agents such as Charlie, the “AI that works for you.” [88]; with most user information stored across personal data stores such as Solid Pods [61, 62, 63].

In this vision, we hypothesise that all data sent from personal data stores, and between personal agents will need to be annotated with terms of use to enable automated compliance with data governance requirements; as data is sent between agents representing data subjects with a range of preferences and legal rights, and hosted in personal data stores across a range of legal jurisdictions. These, and a range of other factors, will influence the terms of use that recipients must comply with when receiving data they wish to process. Specifically, unlike cookies on websites, secondary data transmission (of original and downstream data) and data combination are expected to be more prominent, both in spatial and temporal manners, in this interaction model, which requires special attention in policy languages and engines.

Our hope is that by introducing a ‘Data-Policy’ header where agreed-upon terms of use can be exchanged between clients and browsers; we prove the concept of terms of use agreements at Web scale; and this ‘Data-Policy’ header can be extended and re-used to support HTTP-based data sharing between Web agents, personal data stores and data processors.

4.3. Comparison to Consent Management Platforms

Today, most Websites use Consent Management Platforms (CMPs) to manage their data privacy obligations related to cookies. Consent Management Platforms (CMPs) have 4 primary roles [89]: **Consent collection** via cookie banners; **Consent management** blocking cookies which users have not consented to the use of; **Consent signals** sharing the collected consent with first- and third-party data processors, such as analytics platforms and ad vendors; **Proof of consent** storing proof of consent for regulatory purposes.

For the sharing of consent signals, each third party vendor implements custom consent API’s such as the [Google Consent API](#), and CMPs bear the responsibility of translating the collected user consent into a fixed set of consent options offered by the 3rd party vendor. Our proposed introduction of the ‘Data-Purpose’ header offers a better experience for both vendors and end users. Browsers become responsible for **consent collection** and **consent management**

of cookies; **consent signals** are sent directly from the browser to first and third parties via ‘Data-Purpose’ header, removing the need for intermediary management of consent signals; and **proof of consent** is made possible by simply maintaining a log of the ‘Data-Purpose’ headers that Websites receive. Consequently, our proposal makes the flow of user consent records more rigorous, both by having unambiguous machine-interpretable records of the terms of use that users have agreed to; and having these agreements sent directly to the relevant data processor rather than requiring out-of-band communication via CMPs.

4.4. Comparison to related work on privacy policies in browsers

The core differences between our proposal and the related works have to do with the expressivity of the languages that we propose to use, and the differing legal context at the time in which we do our work.

P3P [64] contains many similar concepts to those which we propose here, including having the ability to define cookie purposes (although from a fixed vocabulary) and a trust engine to mediate between user preferences descriptions of cookie purposes. Regarding expressivity, P3P proposes describing the following features of cookies: (1) *Categories* – what information is collected, (2) *Purpose* – how it is used, (3) *Recipient* – who has access to it, (4) *Retention* – how long it is stored, and (5) *Access* – what information can the user access. ODRL and DToU express a superset of these concepts. P3P only offers 10 purpose categories (and one ‘other’) category that are built into the specification and thus not extensible. In contrast, DPV currently has 95 purposes, and is extensible through OWL [90] – with the ability to preserve semantic relationships. For instance, Websites could define ‘marketingDigitalProducts’ as a subclass of ‘marketing’ to improve precision of their terms of use requests; and browsers would still be able to apply all user preferences for ‘marketing’ preferences to the request.

We recognise that there have been attempts to extend P3P with policies modelled in RDF [91] using Rei [92]. The primary goal of this work by Kolari et al. [91] was to improve expressivity and adoption of P3P. The primary advantage of our proposed use of ODRL or DToU with DPV is (1) ODRL and DPV are more mature than Rei [92], (2) DPV has an extensive range of terms available for describing a wide array of privacy concepts, (3) DToU provides a unified framework covering more concept scopes envisioned in this paper and (4) these vocabularies are built with modern regulation, such as GDPR [57], in mind.

Compared to P3P, our proposal changes the party controlling the terms of use agreement. With P3P, Websites declare the privacy policies associated with different components of the site; similar to the terms of use requests that Websites make in our proposal. However, in P3P the browser is not able to respond with a modified agreement when it sends data, which may choose to allow the Website to use the data for a subset of the purposes it had requested. Instead, in P3P the browser only has the option to block the widget which sends the data; making that functionality completely unavailable to the user.

The **Do Not Track** (DNT) and **Global Privacy Control** (GPC) initiatives are less expressive by design, only offering a single signal to indicate that users do not wish to be tracked via cookies.

In terms of the legal context, there is some level of consensus that P3P and DNT were both unsuccessful due to a lack of legal pressures, however, as evidenced by the greater success

of Global Privacy Control, there is a greater promise of adoption for such technologies when they are mandated in regulation such as the CCPA. This is why we do not propose an isolated technical solution, but rather a collaborative development between research, industry and regulatory bodies to work towards a sociotechnical solution by which the vocabularies used for formally describing terms of use agreements between the user and the website contain terms and concepts that have a well-understood legal interpretation.

4.5. Call to action

4.5.1. Regulators

We call upon legal and policy experts working in the space of data governance to participate in co-designing machine-readable cookie description and transmission standards. In particular, we call upon supervisory bodies, such as the European Data Protection Board (EDPB) or the Information Commissioner’s Office (ICO), that are capable of enforcing compliance with standards such as those we propose for terms of use descriptions, to ensure compatibility between the architectures we build, and the regulatory frameworks of the regions in which they will be deployed.

In the European context, we would like to work with EDPB to understand how these transmitted terms of use can become legally binding Data Sharing Agreements (DSAs), or be considered lawful consent for data processing by data subjects. The ideal outcome of this work-item is a set of terms, and flows, which all EU Data Protection Authorities recognise as valid DSAs or lawful consent under GDPR [57]. In the UK, we would like to work with the ICO to work towards a similar understanding of how transmitted terms of use can constitute legally binding DSAs or lawful consent under the UK Data Protection Act 2018 (c. 12) [93].

A concrete starting point is to establish how terms of use annotations in the ‘Data-Policy’ header sent by browsers can constitute lawful consent for data processing. In particular asking the question: how do we ensure that these terms of use annotations constitute lawful consent when a browser or browser extension computes these terms of use annotations on a users’ behalf?

4.5.2. Industry

At the same time, we call upon industry, including CMP providers, to co-design a solution that will benefit the customer experience and also add value to companies that implement the standard. Secondly, we call upon industry and research centres, such as the Joint Research Centres (JRCs) supported by the European Commission, that have experience implementing mechanised data governance to participate in implementing and validating the proposed policy languages for terms of use exchanges. In particular, we seek to reduce the friction in the adoption of this Web standard by having the formal policy descriptions easily map to internal architectures for enterprise data governance.

5. Conclusion and Future Work

The work we have presented in this paper brings us closer to agentic representation of legal entities at Web Scale.

In Section 3 we implemented a generic personal assistant that communicates using a protocol satisfying the requirements of the Web-scale agentic communication protocol presented in Section 3.1. Future work in this direction will make the design requirements more rigorous by (1) gathering requirements for personal agents through user studies, and (2) engaging with industry to develop specialised agents, including product sales agents. Concurrently, we shall formalise the vocabularies for exchanging *provenance* and *terms of use* between agents and modelling *trust* and *data policies* within agents, extending those vocabularies discussed in Section 3.2. Once these vocabularies mature, we will develop reasoning specifications to mediate between the internal representations and exchanged metadata. This enables agents to negotiate to obtain sufficient provenance to believe claims, and find agreeable data terms of use between agents - whilst concurrently updating their internal models via user interaction.

In Section 4 we looked further into realising semi-automated and legally binding *terms of use* exchanges between agents. We presented a comprehensive vision for the future of automated and transparent data governance on the Web, specifically focusing on the management of cookies. Our proposed solution leverages policy languages such as ODRL, DToU, and DPV to describe cookie policies and data usage agreements in a machine-readable format. This approach aims to enhance user control over their privacy, streamline compliance for businesses, and facilitate regulatory oversight.

These contributions lay the groundwork for digital agents that empower users to fully leverage the benefits of the Web's vast services and information, while protecting them from algorithmic biases, misinformation, and exploitative usage agreements.

References

- [1] M. Wooldridge, *An introduction to multiagent systems*, John Wiley & Sons, 2009.
- [2] M. Wooldridge, N. R. Jennings, *Intelligent agents: Theory and practice*, *The knowledge engineering review* 10 (1995) 115–152.
- [3] A. Olivé, *Conceptual modeling of information systems*, Springer Science & Business Media, 2007.
- [4] R. Verborgh, E. Mannens, R. Van de Walle, The rise of the web for agents, in: *Proceedings of the First International Conference on Building and Exploring Web Based Environments*, 2013, pp. 69–74.
- [5] R. Verborgh, What Web agents want, 2013. URL: <https://ruben.verborgh.org/blog/2013/01/31/what-web-agents-want/>.
- [6] R. Verborgh, Reflections of knowledge, 2021. URL: <https://ruben.verborgh.org/blog/2021/12/23/reflections-of-knowledge/>.
- [7] S. McIlraith, T. Son, H. Zeng, Semantic web services, *IEEE Intelligent Systems* 16 (2001) 46–53. doi:10.1109/5254.920599.
- [8] M. Wooldridge, G. M. O’Hare, R. Elks, *Feline: A case study in the design and implementation of a co-operating expert system* (1991).
- [9] R. Verborgh, J. De Roo, Drawing conclusions from linked data on the web: The eye reasoner, *IEEE Software* 32 (2015) 23–27.
- [10] D. Wood, M. Lanthaler, R. Cyganiak, *RDF 1.1 Concepts and Abstract Syntax*, W3C Recommendation 25 February 2014 (2014). URL: <https://www.w3.org/TR/rdf11-concepts/>.
- [11] OpenAI Platform, ????. URL: <https://platform.openai.com>.
- [12] L.-P. Meyer, C. Stadler, J. Frey, N. Radtke, K. Junghanns, R. Meissner, G. Dziwis, K. Bulert, M. Martin, Llm-assisted knowledge graph engineering: Experiments with chatgpt, in: *Working conference on Artificial Intelligence Development for a Resilient and Sustainable Tomorrow*, Springer Fachmedien Wiesbaden Wiesbaden, 2023, pp. 103–115.
- [13] S. Pramanik, J. Alabi, R. S. Roy, G. Weikum, Uniqorn: unified question answering over rdf knowledge graphs and natural language text, *Journal of Web Semantics* (2024) 100833.
- [14] A. d. Garcez, L. C. Lamb, Neurosymbolic AI: the 3rd wave, *Artificial Intelligence Review* 56 (2023) 12387–12406. URL: <https://doi.org/10.1007/s10462-023-10448-w>. doi:10.1007/s10462-023-10448-w.
- [15] J. S. Evans, In two minds: dual-process accounts of reasoning, *Trends in Cognitive Sciences* 7 (2003) 454–459. URL: <https://doi.org/10.1016/j.tics.2003.08.012>. doi:10.1016/j.tics.2003.08.012, publisher: Elsevier.
- [16] K. Zheng, K. Zhou, J. Gu, Y. Fan, J. Wang, Z. Di, X. He, X. E. Wang, Jarvis: A neuro-symbolic commonsense reasoning framework for conversational embodied agents, 2022. URL: <https://arxiv.org/abs/2208.13266>. arXiv:2208.13266.
- [17] T. Berners-Lee, 30 years on, what’s next# fortheweb? (2019).
- [18] A. Skulmowski, K. M. Xu, Understanding Cognitive Load in Digital and Online Learning: a New Perspective on Extraneous Cognitive Load, *Educational Psychology Review* 34 (2022) 171–196. URL: <https://doi.org/10.1007/s10648-021-09624-7>. doi:10.1007/s10648-021-09624-7.
- [19] V. Vuori, N. Helander, J. Okkonen, Digitalization in knowledge work: the dream of

- enhanced performance, *Cognition, Technology & Work* 21 (2019) 237–252. URL: <https://doi.org/10.1007/s10111-018-0501-3>. doi:10.1007/s10111-018-0501-3.
- [20] F. Brachten, F. Brünker, N. R. J. Frick, B. Ross, S. Stieglitz, On the ability of virtual agents to decrease cognitive load: an experimental study, *Information Systems and e-Business Management* 18 (2020) 187–207. URL: <https://doi.org/10.1007/s10257-020-00471-7>. doi:10.1007/s10257-020-00471-7.
- [21] C. M. de Melo, K. Kim, N. Norouzi, G. Bruder, G. Welch, Reducing Cognitive Load and Improving Warfighter Problem Solving With Intelligent Virtual Assistants, *Frontiers in Psychology* 11 (2020). URL: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2020.554706>. doi:10.3389/fpsyg.2020.554706.
- [22] K. Kim, C. M. de Melo, N. Norouzi, G. Bruder, G. F. Welch, Reducing task load with an embodied intelligent virtual assistant for improved performance in collaborative decision making, in: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2020, pp. 529–538. doi:10.1109/VR46266.2020.00074.
- [23] R. Bregman, *Utopia for realists*, Bloomsbury Publishing, 2018.
- [24] E. Kubin, C. von Sikorski, The role of (social) media in political polarization: a systematic review, *Annals of the International Communication Association* 45 (2021) 188–206. URL: <https://doi.org/10.1080/23808985.2021.1976070>. doi:10.1080/23808985.2021.1976070. arXiv:<https://doi.org/10.1080/23808985.2021.1976070>.
- [25] J. M. Twenge, J. Haidt, T. E. Joiner, W. K. Campbell, Underestimating digital media harm, *Nature Human Behaviour* 4 (2020) 346–348. URL: <https://doi.org/10.1038/s41562-020-0839-4>. doi:10.1038/s41562-020-0839-4.
- [26] T. H. Davenport, J. C. Beck, The attention economy, *Ubiquity* 2001 (2001) 1–es.
- [27] R. Clarke, Risks inherent in the digital surveillance economy: A research agenda, *Journal of information technology* 34 (2019) 59–80.
- [28] D. Searls, *The intention economy: when customers take charge*, Harvard Business Press, 2012.
- [29] Y. Ran, Y. Zeng, Y. Dong, S. X. Zhu, M. Wu, Optimizing pricing and ordering strategies for new products in the presence of consumers with pre-purchase beliefs, *Annals of Operations Research* 337 (2024) 313–342. URL: <https://doi.org/10.1007/s10479-024-05894-w>. doi:10.1007/s10479-024-05894-w.
- [30] P. Whitney, Design and the economy of choice, *She Ji: The Journal of Design, Economics, and Innovation* 1 (2015) 58–80.
- [31] Z. Tufekci, Algorithmic harms beyond facebook and google: Emergent challenges of computational agency, *Colo. Tech. LJ* 13 (2015) 203.
- [32] O. Lassila, J. Hendler, T. Berners-Lee, The semantic web, *Scientific American* 284 (2001) 34–43.
- [33] S. Luke, L. Spector, D. Rager, J. Hendler, Ontology-based web agents, in: *Proceedings of the first international conference on Autonomous agents*, 1997, pp. 59–66.
- [34] S. Poslad, Specifying protocols for multi-agent systems interaction, *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 2 (2007) 15–es.
- [35] N. Shadbolt, T. Berners-Lee, W. Hall, The semantic web revisited, *IEEE Intelligent Systems* 21 (2006) 96–101. doi:10.1109/MIS.2006.62.
- [36] Y. Deng, A. Zhang, Y. Lin, X. Chen, J.-R. Wen, T.-S. Chua, Large language model powered

- agents in the web, learning 2 (2024) 20.
- [37] L. Sun, Y. Huang, H. Wang, S. Wu, Q. Zhang, C. Gao, Y. Huang, W. Lyu, Y. Zhang, X. Li, et al., Trustllm: Trustworthiness in large language models, arXiv preprint arXiv:2401.05561 (2024).
 - [38] H. Story, T. Berners-Lee, A. Samba, R. Taelman, J. Scazzosi, Web Identity (WebID) 1.0, W3C Community Group Final Report, W3C, 2024. <https://w3c.github.io/WebID/spec/identity/>.
 - [39] M. Bosquet, Access Control Policy (ACP), Solid Editor’s Draft, W3C, 2022. <https://w3c.github.io/WebID/spec/identity/>.
 - [40] R. Iannella, S. Villata, ODRL Information Model 2.2, 2023. URL: <https://www.w3.org/TR/odrl-model/>.
 - [41] H. J. Pandit, A. Polleres, B. Bos, R. Brennan, B. Bruegger, F. J. Ekaputra, J. D. Fernández, R. G. Hamed, E. Kiesling, M. Lizar, et al., Creating a vocabulary for data privacy: The first-year report of data privacy vocabularies and controls community group (dpvcg), in: On the Move to Meaningful Internet Systems: OTM 2019 Conferences: Confederated International Conferences: CoopIS, ODBASE, C&TC 2019, Rhodes, Greece, October 21–25, 2019, Proceedings, Springer, 2019, pp. 714–730.
 - [42] J. Wright, B. Esteves, R. Zhao, Me want cookie! towards automated and transparent data governance on the web, 2024. URL: <https://arxiv.org/abs/2408.09071>. arXiv: 2408.09071.
 - [43] M. Garcia, What Air Canada Lost In ‘Remarkable’ Lying AI Chatbot Case, <https://www.forbes.com/sites/marisagarcia/2024/02/19/what-air-canada-lost-in-remarkable-lying-ai-chatbot-case/>, 2024. [Accessed 05-07-2024].
 - [44] V. K. Kommineni, B. König-Ries, S. Samuel, From human experts to machines: An llm supported approach to ontology and knowledge graph construction, 2024. URL: <https://arxiv.org/abs/2403.08345>. arXiv: 2403.08345.
 - [45] T. Berners-Lee, Notation3, <http://www.w3.org/DesignIssues/Notation3.html> (1998).
 - [46] D. Longley, G. Kellogg, D. Yamamoto, M. Sporny, Access Control Policy (ACP), Solid Editor’s Draft, W3C, 2022. <https://w3c.github.io/WebID/spec/identity/>.
 - [47] M. Richardson, R. Agrawal, P. Domingos, Trust management for the semantic web, in: International semantic Web conference, Springer, 2003, pp. 351–368.
 - [48] S. Galizia, Wsto: A classification-based ontology for managing trust in semantic web services, in: European semantic web conference, Springer, 2006, pp. 697–711.
 - [49] W. Sherchan, S. Nepal, J. Hunklinger, A. Bouguettaya, A trust ontology for semantic services, in: 2010 IEEE International Conference on Services Computing, IEEE, 2010, pp. 313–320.
 - [50] G. Amaral, T. P. Sales, G. Guizzardi, D. Porello, Towards a reference ontology of trust, in: On the Move to Meaningful Internet Systems: OTM 2019 Conferences: Confederated International Conferences: CoopIS, ODBASE, C&TC 2019, Rhodes, Greece, October 21–25, 2019, Proceedings, Springer, 2019, pp. 3–21.
 - [51] A. Bouhoula, K. Kubicek, A. Zac, C. Cotrini, D. Basin, Automated large-scale analysis of cookie notice compliance, 2024. URL: <https://www.usenix.org/conference/usenixsecurity24/presentation/bouhoula>, preprint.
 - [52] J. A. Obar, A. Oeldorf-Hirsch, The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services, Information, Communication

- & Society 23 (2020) 128–147.
- [53] R. Iannella, S. Michael, S. Myles, V. Rodríguez-Doncel, ODRL Vocabulary and Expression 2.2, 2018. URL: <https://www.w3.org/TR/odrl-vocab/>.
 - [54] R. Zhao, J. Zhao, Perennial semantic data terms of use for decentralized web, in: Proceedings of the ACM on Web Conference 2024, WWW '24, ACM, 2024. URL: <http://dx.doi.org/10.1145/3589334.3645631>. doi:10.1145/3589334.3645631.
 - [55] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data Privacy Vocabulary (DPV)–Version 2, arXiv preprint arXiv:2404.13426 (2024).
 - [56] B. Adida, et al., RDFa in XHTML: Syntax and Processing, W3C Rec. 14 Oct. 2008, W3C Semantic Web Deployment WG, 2008.
 - [57] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), Official Journal of the European Union L 119 (2016) 1–88. URL: <http://data.europa.eu/eli/reg/2016/679/oj>.
 - [58] Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), 2002. URL: <http://data.europa.eu/eli/dir/2002/58/oj/eng>.
 - [59] California Consumer Privacy Act, 2018. URL: https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375.
 - [60] T. Berners-Lee, J. Hendler, O. Lassila, The semantic web, Scientific American 284 (2001) 34–43.
 - [61] A. V. Sambra, E. Mansour, S. Hawke, M. Zereba, N. Greco, A. Ghanem, D. Zagidulin, A. Abounaga, T. Berners-Lee, Solid: a platform for decentralized social applications based on linked data, MIT CSAIL & Qatar Computing Research Institute, Tech. Rep. (2016).
 - [62] S. Capadisli, T. Berners-Lee, R. Verborgh, K. Kjernsmo, Solid Protocol, 2021. URL: <https://solidproject.org/TR/2021/protocol-20211217>.
 - [63] R. Verborgh, Re-decentralizing the Web, For Good This Time, 1 ed., Association for Computing Machinery, New York, NY, USA, 2023, p. 215–230. URL: <https://doi.org/10.1145/3591366.3591385>.
 - [64] J. Reagle, L. F. Cranor, The platform for privacy preferences, Communications of the ACM 42 (1999) 48–55.
 - [65] J. EPIC, Pretty Poor Privacy: An Assessment of P3P and Internet Privacy, <https://archive.epic.org/reports/pretypoorprivacy.html>, 2000.
 - [66] S. Zimmeck, P. Snyder, J. Brookman, A. Zucker-Scharff, Global Privacy Control – Take Control Of Your Privacy, 2024. URL: <https://globalprivacycontrol.org/>.
 - [67] S. Zimmeck, P. Snyder, J. Brookman, A. Zucker-Scharff, Global Privacy Control (GPC), Proposal 22 March 2024, W3C, 2024. URL: <https://privacycg.github.io/gpc-spec/>.
 - [68] R. T. Fielding, D. Singer, Tracking Preference Expression (DNT), Working Group Note 17 January 2019, W3C, 2019. URL: <https://www.w3.org/TR/tracking-dnt/>.
 - [69] S. L. Pardau, The California Consumer Privacy Act: Towards a European-Style Privacy Regime in the United States, Journal of Technology Law & Policy 23 (2018) 68–114.
 - [70] S. Zimmeck, O. Wang, K. Alicki, J. Wang, S. Eng, Usability and Enforceability of Global

- Privacy Control, *Proceedings on Privacy Enhancing Technologies 2023* (2023) 265–288. doi:10.56553/popets-2023-0052.
- [71] G. Bushati, S. C. Rasmusen, A. Kurteva, A. Vats, P. Nako, A. Fensel, What is in your cookie box? Explaining ingredients of web cookies with knowledge graphs, *Semantic Web* (2023) 1–17. doi:10.3233/SW-233435.
- [72] M. Nouwens, I. Liccardi, M. Veale, D. Karger, L. Kagal, Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence, in: *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–13.
- [73] T. H. Soe, O. E. Nordberg, F. Guribye, M. Slavkovik, Circumvention by design – dark patterns in cookie consent for online news outlets, in: *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society, NordiCHI '20*, Association for Computing Machinery, New York, NY, USA, 2020. doi:10.1145/3419249.3420132.
- [74] G. Kampanos, S. F. Shahandashti, Accept All: The Landscape of Cookie Banners in Greece and the UK, in: A. Jøsang, L. Fitcher, J. Hagen (Eds.), *ICT Systems Security and Privacy Protection*, Springer International Publishing, Cham, 2021, pp. 213–227.
- [75] H. Habib, S. Pearman, J. Wang, Y. Zou, A. Acquisti, L. F. Cranor, N. Sadeh, F. Schaub, "It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 1–12. doi:10.1145/3313831.3376511.
- [76] C. Santos, A. Rossi, L. Sanchez Chamorro, K. Bongard-Blanchy, R. Abu-Salma, Cookie banners, what's the purpose? Analyzing cookie banner text through a legal lens, in: *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 2021, pp. 187–194.
- [77] D. Bollinger, K. Kubicek, C. Cotrini, D. Basin, Automating cookie consent and GDPR violation detection, in: *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2893–2910.
- [78] B. Adida, M. Birbeck, S. McCarron, I. Herman, RDFa Core 1.1, W3C Recommendation 17 March 2015, 2015. URL: <https://www.w3.org/TR/rdfa-core/>.
- [79] R. T. Fielding, M. Nottingham, J. Reschke, HTTP Semantics, RFC 9110, 2022. doi:10.17487/RFC9110.
- [80] M. Degeling, C. Utz, C. Lentzsch, H. Hosseini, F. Schaub, T. Holz, We Value Your Privacy... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy, *CoRR abs/1808.05096* (2018). URL: <http://arxiv.org/abs/1808.05096>.
- [81] L. F. Cranor, Necessary but not sufficient: Standardized mechanisms for privacy notice and choice, *J. on Telecomm. & High Tech. L.* 10 (2012) 273.
- [82] P. Ryan, M. Crane, R. Brennan, GDPR Compliance tools: best practice from RegTech, in: *International Conference on Enterprise Information Systems*, Springer, 2020, pp. 905–929.
- [83] O. Olawale, F. A. Ajayi, C. A. Udeh, O. A. Odejide, RegTech innovations streamlining compliance, reducing costs in the financial sector, *GSC Advanced Research and Reviews* 19 (2024) 114–131.
- [84] V. Morel, C. Santos, V. Fredholm, A. Thunberg, Legitimate Interest is the New Consent – Large-Scale Measurement and Legal Compliance of IAB Europe TCF Paywalls, in:

- Proceedings of the 22nd Workshop on Privacy in the Electronic Society, CCS '23, ACM, 2023. doi:10.1145/3603216.3624966.
- [85] J. H. Schumann, F. Von Wangenheim, N. Groene, Targeted online advertising: Using reciprocity appeals to increase acceptance among users of free web services, *Journal of Marketing* 78 (2014) 59–75.
 - [86] D. J. Solove, Murky consent: an approach to the fictions of consent in privacy law, *BUL Rev.* 104 (2024) 593.
 - [87] M. Florea, B. Esteves, Is Automated Consent in Solid GDPR-Compliant? An Approach for Obtaining Valid Consent with the Solid Protocol, *Information* 14 (2023). doi:10.3390/info14120631.
 - [88] Tim Berners-Lee, Charlie works for Bob, Technical Report, w3c, 2024. <https://www.w3.org/DesignIssues/Charlie.html>.
 - [89] C. Team, Consent Management Platform (CMP): How Does it Work?, 2024. URL: <https://www.cookieyes.com/blog/consent-management-platform/>.
 - [90] B. Motik, P. F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler, et al., OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, W3C Recommendation 11 December 2012, 2012.
 - [91] P. Kolari, L. Ding, L. Kagal, S. Ganjugunte, A. Joshi, T. Finin, et al., Enhancing P3P framework through policies and trust, UMBC Technical Report, TR-CS-04-13, 2004.
 - [92] L. Kagal, T. Berners-Lee, Rein: Where policies meet rules in the semantic web, *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA* 2139 (2005).
 - [93] Data protection act 2018 (c. 12), 2018. URL: <https://www.legislation.gov.uk/ukpga/2018/12/contents/enacted>.