Combating the Effects of Cyber-Psychosis: Using Object Security to Facilitate Critical Thinking

Robert H. Thomson¹, Quan Nguyen², Essien Ayanam², Matthew Canham³, Thomas C. Schmidt⁴, Matthias

Wählisch⁵ & Eric Osterweil²

¹Army Cyber Institute, U.S. Military Academy, West Point, USA, robert.thomson@westpoint.edu

²Department of Computer Science, George Mason University, Fairfax, USA,

{qnguye31,eayanam,eoster}@gmu.edu

³Cognitive Security Institute, matthew@cognitivesecurity.email

⁴HAW Hamburg, Hamburg, Germany, t.schmidt@haw-hamburg.de

⁵TU Dresden, Dresden, Germany, m.waehlisch@tu-dresden.de

Abstract—Humanity is currently facing an existential crisis about the nature of truth and reality driven by the availability of information online which overloads and overwhelms our cognitive capabilities, which we call Cyber-Psychosis. The results of this Cyber-Psychosis include the decline of critical thinking coupled with deceptive influences on the Internet which have become so prolific that they are challenging our ability to form a shared understanding of reality in either the digital or physical world. Fundamental to mending our fractured digital universe is establishing the ability to know where a digital object (i.e., a piece of information like text, audio, or video) came from, whether it was modified, what it is derived from, where it has been circulated, and what (if any) lifetime that information should have. Furthermore, we argue that on-by-default object security for genuine objects will provide the necessary grounding to support critical thinking and rational online behavior, even with the ubiquity of deceptive content. To this end, we propose that the Internet needs an object security service layer. This proposition may not be as distant as it may first seem. Through an examination of several venerable (and new) protocols, we show how pieces of this problem have already been addressed. While interdisciplinary research will be key to properly crafting the architectural changes needed, here we propose an approach for how we can already use fallow protections to begin turning the tide of this emerging Cyber-Psychosis today!

Index Terms—cognitive security, data security, decisionmaking, psychology, object security, protocols

I. INTRODUCTION

I N the last decade, the average adult makes upwards of 35,000 decisions per day [1], up from an estimate of around 9,000 during working hours in 1994 [2]. Among key differences in the last 30 years are the *sources* and fidelity of information and evidence. This shift has been fueled by the rise of the Internet, search engines, social media, and online shopping. Exponentially more (lower quality) information is readily available to anyone with an Internet connection. While initially seen as a beneficial source for the democratization of information [3], this notion has been challenged by the rise

of *deceptive content* [4], algorithmic feeding of news creating *filter bubbles* [5] causing echo chambers [6] and polarization [7], and the paywalling of reliable academic content leading to the relatively greater availability of less rigorous research online [8].

At the same time, there is evidence that increased use of the Internet and search engines - also known as the Google Effect [9] which is also being exacerbated by recent use of generative artificial intelligence (genAI) in education [10], has led to a reduction in critical thinking and a failure of metacognitive calibration (e.g., an overconfidence in one's depth of understanding of a given topic; [11]). The sheer scale of information available means that it is possible to find a contrasting opinion on nearly any topic. The information we consume almost exclusively takes form of "digital objects." These include images, news stories, videos, audio recordings, and more. They have their own lifecycles, origins, and targeting algorithms. While many people believe that humans forage for information through a sensemaking process [12], in the Information Age it is information which is algorithmically targeted to us [5]. In a very real sense, data in digital objects should be reasoned about as if it has its own motivations [13], or its own "agency." When one additionally factors in the opacity of the algorithms which feed us information, it becomes nearly impossible to judge the veracity of a claim. As such, we argue that the current information space leads to "Cyber-Psychosis," which will only become more problematic as we increasingly base our real-life decisions on information in digital objects and as technology pushes us into the metaverse [14].

In this paper, we argue that there *is* a viable and immediate path toward restoring people's ability to discern fact from fiction online. We illustrate that *critical thinking* skills together with *object security* can combat the roots of Cyber-Psychosis. Further, basic object security primitives have existed in the core of the Internet for decades, in protocols ranging from Secure/Multipurpose Internet Mail Extensions (S/MIME) [15], to the Resource Public Key Infrastructure (RPKI) [16], and the more recent Coalition for Content Provenance and Authenticity (C2PA) [17]. Moreover, a combination of *existing* generalby-design cybersecurity protocols and protections, such as the Domain Name System's (DNS') [18] Security Extensions

This paper was funded in part from an Office of Naval Research FY21 Multi-University Research Award (N0001423MP00219 & N0001425MP00064) to RT.

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

(DNSSEC) [19]–[22] and the DNS-based Authentication of Named Entities (DANE) [23]–[26] are poised and ready to be used and built upon. As such, we issue a **call-to-arms** to develop interdisciplinary frameworks to help restore the security and resilience of our cognitive processes from deceptive online content. We argue that a first step includes evolving Internet services to produce secure-by-default digital objects via a generalized *object security service layer*. This strategic goal will require basic research and engineering, but we also show that many of the most basic necessary components for us to address our tactical needs are already deployed and can make a meaningful difference if used today.

In the remainder of this paper, we begin by operationalizing the term Cyber-Psychosis and show how the present information space parallels the effects of clinical psychosis in Section II. We then proceed, in Section III, to evaluate critical thinking frameworks used in education as a way to use reasoning strategies to ground out decision-making processes, and show a common set of principles which would mitigate the effect of Cyber-Psychosis. Using these critical thinking frameworks for guidance, we present a mapping between what they need, and how we can be (and in some cases already are) able to secure these with Internet object security in Section IV. Next, in Section V, we propose a structured research agenda to bridge the gap between the existing protocol-specific object security service sets and research needed to realize a generalized object security service layer for the Internet. Then, we detail several prime example protocols which have already operationalized important aspects of generalized object security in Section VI. Finally, we present a discussion of the state of our Internet in Section VII and then conclude with our call-to-arms in Section VIII.

II. DEFINING CYBER-PSYCHOSIS

In psychiatry, the hallmark of a psychotic episode is the loss of one's sense of reality [27]. Cyber-Psychosis has been a term often used in science fiction to describe a mental condition that arises from excessive interaction with our dependence on technology, and often results in some kind of dehumanization. Originally popularized by William Gibson's *Neuromancer* [28] and more recently by the video game CyberPunk 2077 [29], the concept involves people losing touch with reality due to their immersion in digital environments or cybernetic enhancements. While originally being an allegory for dehumanization and loss of autonomy by technology [30], we argue that Cyber-Psychosis should actually be defined in line with the clinical definition for psychosis, which includes:

- 1) **Delusions:** Strongly held false beliefs that are resistant to reason or contrary evidence. In the extreme case, these can include paranoia and grandiose thoughts.
- 2) **Hallucinations:** Sensory experiences that appear real but are created by the mind. The most common type is auditory hallucinations, such as hearing voices that others do not hear.
- Disorganized Thinking: This can manifest as incoherent speech, difficulty organizing thoughts, or jumping from one topic to another without logical connection.

 Negative Symptoms: These include diminished emotional expression, lack of motivation, and withdrawal from social activities.

To best address these parallels, we will discuss each feature of psychosis in turn.

A. Delusions

A common discussion in influence and media is the debate surrounding polarization in online spaces. Polarized thought reflects a strongly-held belief that is resistant to counterarguments and contrary evidence. In fact, the contested *backfire effect* [31]–[33] occurs when a person's belief in the polarized topic actually strengthens when presented with sound contradictory evidence. Furthermore, the *continued influence effect* [34]–[36] occurs when discredited prior beliefs continue to influence present thought, which has become prevalent in the current age of misinformation. These distortions in rational thought may be exacerbated by the Google effect insofar as people are overconfident in their understanding of a given topic, and there will almost always be some digital object containing a piece of confirming evidence on the Internet from which to anchor their prior beliefs.

B. Hallucinations

While the hallmark of hallucination is a sensory experience which appears real but isn't, we argue that a metaphorical hallucination occurs during the sensemaking process when evaluating information on the Internet. While it seems like it is the user who is searching for information, the underlying platform's algorithms are selectively feeding data back to the user. These algorithms are often opaque to the user, and this underlying agency of the algorithm is applied to the digital objects whose data may serve to steer the user to the algorithm's own aims (that of the company which made the algorithm) [13]. The focus of social media companies is to feed more ads to generate more revenue, thus the primary goal of information is not to satisfy the users request, but instead to maximize the user's continued interaction with the platform to generate more revenue. These algorithms effectively act as a hidden voice steering the user to a goal not of the user's own volition. This may only get worse with the advent of the metaverse and further integration of augmented and virtual reality into our daily lives.

C. Disorganized Thinking

In the present discussion, disorganized thinking somewhat overlaps with our other analogies between clinic psychoses and the symptoms of Cyber-Psychosis. The hallmark of disorganized thinking is a difficulty in organizing thoughts and jumping between contexts without an obvious logical connection. For Cyber-Psychosis, this could be a result from the burnout and adrenal fatigue associated with Western society's *always-on* work culture, leading to maladaptive behaviors such as doom-scrolling [37], [38] and outrage fatigue [39]. While not necessarily rising to the level of clinical significance in isolation, it exacerbates the other three features of Cyber-Psychosis discussed in this paper, and this disordered thought and failure of metacognition [11] makes it challenging for the user to detect this impairment.

D. Negative Symptoms

Perhaps the hallmark example of social and emotional withdrawal including a lack of motivation, comes from Hikikomori syndrome [40]. This syndrome, traditionally associated with Japanese culture but is seen around the world, reflects an extreme social isolation, emotional distress, lack of motivation, and addiction to digital spaces. This syndrome has not been officially classified as a mental disorder, but has been associated with prodromal psychosis, the sub-clinical phase predating a full psychotic illness [41]. A challenge with this syndrome is that there are limited effective evidence-based treatments [42]. Most relevant to the current discussion is the role that digital spaces often have in the expression (and likely contributing factor to) this disorder. In the extreme case, the digital world becomes the person's reality. In such cases, being able to quantitatively evaluate the veracity, origins, and provenance of the digital objects that shape a person's world becomes a critical necessity.

E. Consequences of Cyber-Psychosis

We argue that the Internet as currently designed and administered leads to the aforementioned four consequences as seen in numerous bodies of literature, which maps quite well onto the clinic definition for psychosis; thus we describe the current Internet age as leading to Cyber-Psychosis. This not only has a social component, there is evidence that extensive use of digital spaces and technology reduces the level of dopamine [43] and interferes with the effectiveness of the limbic system [44]. This goes beyond pure Internet addiction as a compulsive behavior, it also impacts how we process reality itself. Given the availability of online information that goes well beyond our brain's ability to process [45], this information overload can make it difficult to assess the veracity of information, and we argue object security is a necessary first step. There is evidence that Russia has weaponized the information overload and outrage fatigue facets of Cyber-Psychosis to interfere in foreign elections [46].

Given critical thinking frameworks were developed to assess the veracity of information, how well do they work with digital objects on today's Internet to mitigate the effects of Cyber-Psychosis? Users who have high literacy in algorithmic sources (e.g., chatbots) are more likely to trust their output and use them in their workflow [47], however this use may lead to over-trust [48], [49] and a reduction in critical thinking [10]. While some inoculation (pre-bunking) strategies have proven effective in limiting the acceptance of misinformation specifically [50], we argue that this is insufficient to address the underlying issues of overload and fatigue which plague online spaces. We need object security to allow users to reason about information as *it* encounters and is steered to them.

III. CRITICAL THINKING FRAMEWORKS

A review of recent frameworks endorsed by teachers and librarians (see Table I) highlights a focus on critical reasoning, and it highlights three commonalities: 1) Review the source's authority, expertise, and motivations; 2) review the accuracy and soundness of the argument itself; and 3) understand the motivation behind why the argument exists. Perhaps the most well-studied is CRAAP, which was used extensively in the pre-Internet era to review primarily scientific claims. A concern with this methodology is that it focused primarily on the article itself without reviewing other articles. More recent examples (e.g., CCOW) focus on situating the article in a broader world view, which mitigates some of these concerns and presents the opportunity to understand the context and motivations behind why this information was available to users in the first place [51].

These materials serve an important purpose by highlighting factors for users to focus on when critically assessing a given argument. However, there are some concerns about these processes which may make them more counter-productive than would be seen at first glance. Namely, the focus on source authority and trustworthiness as an early indicator in most frameworks - over the logical form of the argument may lead users to fall prey to source authority biases (e.g., appeals to authority, homophily-induced re-appraisals). For instance, Schoenherr and Thomson [52] (see also [11]) found that identifying the source of an argument as a scientist vs a layperson determined how effective participants were at evaluating the validity of an argument (i.e., its structure having a proper logical form). This went beyond belief bias; the fact that we tend to judge the conclusion of an argument based on how well it fits with our prior beliefs instead of its logical structure. Even more troubling, the study found that participants' confidence scores were relatively unchanged when accuracy fell nearly 20% when inducing belief bias, indicating that there was no dissonance or other metacognitive awareness that their critical thinking was impaired. Soprano et al [53] further identified 39 cognitive biases that may influence fact-checkers in their critical thinking processes. Perhaps a better focus would be to relatively overweight techniques to evaluate the sources of the facts, not the source of the argument itself [54]?

Another concern is that there is reasonable disagreement between methods. Portillo et al [60] found disagreement between critical thinking frameworks' (DISCERN, CRAAP, and HONcode) evaluation of common medical websites including Medline, Healthline, Mayo Clinic, and WebMD. Specifically Medline Plus was the only website which did not receive HONcode certification, while having the highest reliability by DISCERN and CRAAP.

A. Assessable Features of an Argument to Support Critical Thinking

- Source reputation of publication or website. Source reputation of the author.
- Data referenced is accurate. Are articles pointing to source material? Is evidence from reputable source?

| CRAAP | CCOW | RADAR | ACLR | BEAM | | | | |
|----------------------|-----------------------|----------------------|----------------------|----------------------|--|--|--|--|
| Currency: The | Credibility: The | Relevance: The im- | Accuracy: The cor- | Background: The | | | | |
| timeliness of the | trustworthiness of | portance of the in- | rectness of the in- | context or history | | | | |
| information. | the source. | formation for your | formation. | behind the source. | | | | |
| | | needs. | | | | | | |
| Relevance: The im- | Consistency: | Authority: The | Clarity: The clear- | Exhibit: Evidence | | | | |
| portance of the in- | Whether the | qualifications or | ness and ease of un- | or examples | | | | |
| formation for your | information is | expertise of the | derstanding of the | presented by | | | | |
| research. | consistent with | source. | argument. | the source. | | | | |
| | other sources. | | | | | | | |
| Authority: The | Objectivity: The | Date: The publica- | Logic: The reason- | Argument: The | | | | |
| qualifications and | lack of bias or | tion or last updated | ing and structure of | main claim or | | | | |
| credibility of the | impartiality in the | date. | the argument. | thesis the source is | | | | |
| author. | source. | | | trying to prove. | | | | |
| Accuracy: The cor- | Worldview: Critical | Appearance: | Relevance: The | Method: The ap- | | | | |
| rectness of the con- | questioning of the | The professional | importance of the | proach or methodol- | | | | |
| tent and the fact- | source and its con- | presentation and | information in the | ogy used in the re- | | | | |
| checking process. | text as well as one's | design of the | given context. | search. | | | | |
| | own views. | source. | | | | | | |
| Purpose: The rea- | | Reason: The logic | | | | | | |
| son the information | | behind the conclu- | | | | | | |
| exists and the in- | | sions or claims. | | | | | | |
| tended audience. | | | | | | | | |
| TABLE I | | | | | | | | |

A LIST OF COMMON CRITICAL THINKING FRAMEWORKS. OTHERS NOT SHOWN HERE INCLUDE DISCERN [55], HONCODE [56], AND SIFT [57], AMONG OTHERS. CONTENT DERIVED FROM [51], [57]–[59]. IT IS A POSITIVE THAT CCOW HIGHLIGHTS OBJECTIVITY AND WORLDVIEW OF THE ARTICLE AND READER; AS WELL AS RADAR AND ACLR HIGHLIGHTING THE LOGICAL STRUCTURE OF THE ARGUMENT.

- Logical consistency (valid and sound points being made).
- Bias/Objectivity. Why is what being said being said. Who else is saying this. Who is saying the counter? Can there be a dashboard of opinions on a topic?
- Clarity and Precision. Is there language used which could (intentionally) lead to obscurity/vagueness
- Relevance: are examples relevant or are they meant to be evokative and to mislead towards affective red herrings and unrelated bits.
- Does it accept and refute a counter-claim (debating 101).

By evaluating the commonalities in these frameworks, we propose that cybersecurity protections can be applied to digital objects in order to allow Internet users to best understand the common features of an argument which support rational human decision-making and which may help mitigate the effects of Cyber-Psychosis.

IV. COMBATING CYBER-PSYCHOSIS WITH OBJECT Security

While modern Internet usage has shifted almost exclusively to *digital objects*, these objects present a foundational challenge. This is because existing cybersecurity protocols are most commonly only designed to protect data for the short windows of time while it is being transmitted, *i.e.*, "*transmission security*." The most common example of this is the Transport Layer Security (TLS) [61] protocol. However, one central security challenge posed by digital objects is that they exist beyond just when they are being transmitted. They exist on various servers and Internet services, are shared between users, and are promoted by platforms' targeting algorithms. Therefore, they must *also* be secured while "at rest." As a result, digital objects require a different type of security, which can provide more protracted protections, which we call "*object security*." This frames a basic challenge of what protections we need for digital objects so as to best protect them in ways that will let users ingest them for critical thinking.

To know whether our objects are (or can be) secure, we must first understand precisely what protections are needed. Using the critical thinking frameworks from Section III as representative examples, inspection of their objective requirements reveals several recurring objective needs. Of the example set described in Table I, at least three of the five examples require origin authentication protections (i.e., being able to securely verify the sources of information). Specifically, origin authentication would be necessary for CRAAP to establish "Authority." Similarly, CCOW would need it for "Credibility," "Objectivity," and "World View." RADAR would also need it for "Authority." In addition, to implement these frameworks on today's Internet, four frameworks would also need combinations of integrity protections and process provenance [62]. Lastly, at least two rely on forms of data lifecycle management. As origin authentication, provenance, and lifecycle emerge as common requirements, protocol-specific examples can be seen to exist in several prominent Internet protocols (detailed in Section VI). However, what is also needed is a basic understanding of what a generalized object security service layer for the Internet should look like. That future definition looms large as an open research challenge.

V. WHAT DOES IT MEAN TO "SECURE" OBJECTS?

Understanding what is needed for generalized object security between separate administrative domains (*e.g.*, competing companies, collaborating universities, different nations, etc.) must be decomposed into basic pieces. For this, we propose a set of research questions to start:

Research question (I), what are the basic nature and requirements of objects' *origin authentication* protections? This challenge will involve locating origins online and being able to securely verify that genuine objects have come from their reported sources. To illustrate, consider an online new story that is shared multiple times, over multiple social media platforms. Recipients should want to securely verify who authored it, i.e., its origin. Further, that origin may not be useful for the recipient unless it is a name that she/he recognizes. To accomplish this with public-key cryptography, object recipients will need to securely learn the public keys of object originators. However, issues ranging from private key compromise to standard cryptographic hygiene dictate that keys will inevitably (if not periodically) need to be changed over time, *i.e.*, have a *lifecycle*. When an origin needs to change its key it should not necessarily need to reissue all of its objects that are verified by that key. As a result of this, the cryptographic keys, themselves, will not be sufficient as identifiers. Identifying and locating inter-administrative data origins and mapping them to changing identifiers on the Internet has frequently been accomplished through unambiguous naming. In the context of origin authentication, these names must be semantically able to unambiguously locate public keys for identified object sources. Research challenges include understanding what the necessary and sufficient natures of a usable namespace for digital objects.

Research question (II), what semantics will the *lifetimes* of objects' protections need? Developing protection-focused guidance for *object lifetime* will involve creation of evaluation frameworks to inform decisions about how long objects should be protected. For example, some digital objects' content may only be intended to be valid for certain predictable periods of time. Therefore, should objects *only* be verifiable in specific time frames? By contrast, some objects' content may be intended to be valid indefinitely. Should those objects have open ended verifiability? If so, should their validities be periodically renewed, or semantically specified as indefinite? During validity periods of objects' lifetimes, their data must have enforceable integrity protections.

The above research question leads directly into question (III), what does an effective relationship between objects' lifetimes and the lifecycles of the cryptographic keys that protect and verify them look like? While an object needs to be validated, it must have a corresponding cryptographic key to do that verification and integrity protection. This means keys cannot expire during periods of time that the objects they cover are expected to be verified by them. Instead, keys must exist and be valid for the periods that objects need them, which forms an implicit requirement that their lifecycles be congruent. Among the immediate research challenges this poses is how to manage situations in which a source's key needs to change (e.g., emergency, planned retirement, etc.)? In situations like these, having objects' keys be located by names will facilitate key transitions [63], but that alone will not guarantee continuity of protections.

Research question (IV), what aspects of *integrity* and *provenance* will be needed? "Provenance" is a word whose common usage often includes application to the concept of "chain-of-custody" (*e.g.*, what is the ownership history of a piece of art). The word has also been a topic of a great deal of the computer science literature. In the past, it has been used to

document audit trails for e-science [62], [64]-[67], and more recently a growing body of research has considered its applicability to combat deepfake images [68]-[71]. In the more recent focus, much attention has been paid to documenting when an object has been altered (such as when images are touchedup or filtered). In these cases, it must be possible to securely inspect and evaluate these changes. We, therefore, cast these protections as a superset of *integrity protections*, which are inherently offered by this type of object provenance. However, while this type of provenance is necessary, it is not sufficient for combating misinformation and disinformation. We contend that it is also necessary to be able to verifiably express the "process(es)" by which objects have been *derived*, or *process* provenance [62]. Said differently, if an object represents a conclusion that is derived from other objects, it is necessary to securely reference these other objects directly. This will likely entail the same origin naming requirement for each component of an object's provenance. Further, with digital objects having their own agency on the Internet, knowing where they have appeared can be as important as what they say, or where they originated (*i.e.*, chain of custody provenance).

Research question (V): will *confidentiality* **protections need different approaches?** Whereas transmission security protocols have extensively protected confidentiality, are their models applicable to generalized object security? As we discuss in Section VI, there are a number of existing protocols that have attempted aspects of an overall object security service layer. While these examples illustrate applicable lessons for our work here, vanishingly few have attempted longitudinal confidentiality protections between separate administrative entities. Does this imply that there are non-obvious complexities in designing confidentiality protections for inter-administrative object security, or perhaps that they are not necessary?

Research question (VI), *evaluability*, how can object creators and recipients *know* if protections are remaining effective, or not? There needs to be a way for content authors, object recipients, and anyone else to *know* if object protections have failed and/or if they have worked. We believe that object security, in particular, will require a robust approach to knowing the protection status of objects. This is, in part, because of the protracted periods of their protections. For example, knowing that an object has valid cryptographic signatures at one point in time does not necessarily tell a relying party that it was protected yesterday, or will be tomorrow. Moreover, if an object needs to remain secure for any critical function, being able to prove that it has or has not been secure will be necessary to deriving trust from it.

This set of research tasks is an example of how the complex nature of object security could be decomposed. Other valid decompositions could lead to different research tasks, but we believe the above set is an important example starting point. However, there exist examples of inter-administrative security protocols on the Internet that have operationalized aspects of these, which both illustrate existence proofs that these protections are attainable, and in some cases serve as *usable* options for generalized object security *today*. We examine these in detail next, in Section VI.

TABLE IIImplementation of object security protections that are available (\checkmark), partially available ((\checkmark)), or not available (\bigstar)

| Protocol | Origin Authentication | Lifetime | Keys-to-Objects | Integrity | Provenance | Confidentiality | Evaluation |
|--------------|-----------------------|----------|-----------------|-----------|------------|-----------------|------------|
| C2PA | × | × | × | 1 | (✔) | × | × |
| S/MIME & PGP | × | × | × | 1 | × | ✓ | × |
| DNSSEC | 1 | 1 | × | 1 | × | × | × |
| RPKI | 1 | 1 | × | 1 | × | × | × |
| Authenticode | × | 1 | 1 | 1 | × | × | × |



Fig. 1. C2PA provenance and integrity is implemented by Manifests.

VI. WHERE HAVE DIGITAL OBJECTS BEEN SECURED?

Some notable security protocols are designed to protect objects, such as the Coalition for Content Provenance and Authenticity (C2PA) [17], the Secure/Multipurpose Internet Mail Extensions (S/MIME) [15], Pretty Good Privacy (PGP) [72], the Domain Name System's (DNS') [18] Security Extensions (DNSSEC) [19]–[21], [63], the Resource Public Key Infrastructure (RPKI) [16], and Windows' Authenticode code signing protocol [73]. These examples range from a couple of years old to decades and their deployments offer insights into generalizable object security. In this Section, we use our candidate research areas (I) to (VI) as guideposts, and Table II summarizes discussions of them.

C2PA: In recent years the C2PA has emerged as a leading organization in the provenance space [68]-[71], *i.e.*, protecting objects like JPEG, PNG, and MPEG. Its protections are focused on describing how media objects have been created and/or modified, and is implemented by a set of constructs called "manifests" (illustrated in Figure 1). Each manifest contains three to four data items: assertions, claims, a signature, and optionally W3C Verifiable Credentials (VCs) [74]. Assertions describe when an image object was taken, what camera took it, what editing tool may have modified it, and generally what action may have been taken on that object. Claims bundle assertions together by a "claim generator" via cryptographic hashes. They indicate who generated the claim, what software was used to generate it, etc. Finally, the claim signature is generated over the whole manifest to vouch for its integrity. Every modification to a media object should result in a new manifest. The total set is accumulated into the object's "manifest store," (from the "origin manifest" to the "active manifest"). The C2PA specification states that trust in objects should be determined by the identity of the signer.

The C2PA specification partially implements origin authentication (our research area (I)) through its use of X.509 certificates. Verification of objects is accomplished by verifying claim signatures using their corresponding X.509 certificate. However, verification of that certificate hinges on its Subject Alternative Name (SAN), which is most commonly a DNS domain name. This inherently forms a basic dependency on the DNS namespace, but these names are not currently used to authenticate origins' keys. C2PA media objects do not offer semantics for lifetime protections (our research area (II)). Any lifetime restrictions are (at best) implicit from certificate validity lifetimes. This specification does not have an explicit relationship defined between media objects and the keys that protect them. Thus, the key to object management (our research area (III)) is not protected by C2PA. Provenance is implemented using X.509 certificates, Concise Binary Object Representation (CBOR), and constructs called (JPEG Universal Metadata Box Formats) JUMBFs. However, this is only one component of what we consider the process provenance in our vision for research area (IV). The provenance of C2PA media objects describes their modification history provenance, and includes integrity protections built in. We propose that this type of modification history of digital objects is necessary, but not sufficient for problems like misinformation and disinformation. This is why we propose the need for process provenance, i.e., which external objects may have been involved (or led to) the creation of objects. While images may not actually be the beneficiaries of this type of protection, they can be included in the process provenance for others (such as news stories derived from them). The C2PA specification does not provide confidentiality (our research area (V)), however recent research from the literature proposes an approach to provide this [69]. Lastly, C2PA's protections do not include an evaluability approach (our research area (VI)).

S/MIME & PGP: These protocols are among the oldest object focused inter-administrative security protections in use today, though their deployment traction has lagged. They are very similar in the protections they offer. Among the primary benefits derived from using these protocols today is their confidentiality protections (our research area (V)). These protocols both protect general object types, anything that can be sent in an email message. However, they do not address the rest of our research areas. Neither protocol implements origin authentication (our research area (I)). These protocols are not able to verify that keys are securely associated with the names they report. They each rely on external mechanisms (X.509 for S/MIME and the Web of Trust for PGP), but neither serves as



Fig. 2. DNSSEC's hierarchical namespace enables secure key learning.

a reliable inter-administrative origin authentication solution, which we describe in detail in previous work [26]. Neither of them implement lifetime protections (area (II)). Neither have semantics or mechanisms to manage the lifecycles of keys to their protected mail objects (research item (III)). They both have integrity protections, but do not have provenance mechanisms (area (IV)). Lastly, neither have an evaluability approach (area (VI)).

DNSSEC: DNSSEC's namespace is general (the DNS is the Internet's de facto name mapping system), but its object type is DNS-specific (as opposed to being generally useful for other applications, like S/MIME and PGP). In DNSSEC, zone names map to public keys (DNSKEY records), and these keys are used to verify signatures attached to DNSSEC's canonical object type: the "RRset," illustrated in Figure 2. Zone administrators generate signatures (in RRSIG records) that cover objects. Research has shown that its deployment size has been growing steadily since its standardization [63], [75]. Measurements also show that the size of DNSSEC validator traffic (*i.e.*, clients) now exceeds 33% of global DNS [76]. Research has also demonstrated the robustness of the policy framework around the root zone and its key management [77], [78].

DNSSEC implements origin authentication (area (I)). Zones deploy their own named keys, zones also aggregate named objects. These names exist in DNS' collision-free domain namespace. In DNSSEC, the namespace itself binds keys to objects. Though the DNS has been ubiquitous for decades, the ongoing growth and evolution of its ecosystem has exposed security challenges that relate to name collisions. Past research [79] has illustrated that the introduction of new generic Top-Level Domains (gTLD) enabled Man-in-the-Middle (MitM) attacks because of name collisions. Subsequent research [80] illustrated remediative actions for this that an object security service layer will need to incorporate. DNSSEC's trials and tribulations over the complexity of securing a namespace should be advisory for alternate namespace proposals. DNSSEC implements lifetime protections (area (II)) through its signatures over objects. Each signature specifies a definitive inception and expiration time. The implication of this approach is that if any given object needs to be verifiable after its previously specified expiration, it will need to be re-signed with a new inception/expiration. Conversely, if an object becomes invalid while it is still within its specific lifetime, there is no recourse for a zone

administrator. The object can be removed from an operational zone, but will still be replayable to validating resolvers and caches. There exists no research literature studying the impact of object lifetimes in DNSSEC. It does implement integrity protections over its objects (through its covering RRSIGs), but it does not implement provenance protections (area (III)). Research area (IV), lifecycle alignment between keys and objects, has begun to be be studied in the literature. In [63] we found that there is a large disconnect between the prescribed processes for key management and the actual operationally deployed processes. One broader lesson is that protocols need to be designed to maintain continuity of protections for covered objects. Research area (V) (confidentiality) is not implemented by DNSSEC. Though, there has been a push for protocols called "encrypted DNS," such as DNS over HTTPS (DoH), DNS over TLS (DoT), etc. These, however, implement transmission security and do not address object security protections. Research addressing evaluation of DNSSEC (area (VI)) has primarily focused on its availability, rather than its object protections. For example, a challenge that was identified by research was a reduced incidence of validation by resolvers (clients) [81], [82]. This research discussed the success rates of validation, but did not examine the success rates of validating the signed objects. This is not a critique of the research and does not reflect a limitation of it. Rather, this is an observation that this research's lessons are not focused on the *object security* implications, but instead the overall deployment's performance.

RPKI: This protocol secures the global routable Internet Protocol (IP) address space and the Autonomous System Number (ASN) space. To compare this with the DNS, names are IP prefixes and the canonical object type is called a Route Origin Attestation (ROA), each of which embeds an End Entity (EE) certificate. The "origin" in ROAs refers to a BGP origin ASN, not the creator of an object. RPKI defines a BGP-level verification process called Route Origin Validation (ROV) [83] to verify the correctness of routing information in ROAs. The RPKI has been undergoing deployment for over 10 years [84], and recent research shows its adoption and deployment traction have been growing [85] since the beginning [86], [87].

The RPKI implements origin authentication (area (I)) by creating a hierarchy of digital certificates that chain from a Regional Internet Registry (RIR) to any ROA. This means that any rightful resource holder of an IP prefix can create a ROA to provably attest to a mapping of that prefix to an ASN. Complementing this, any relying party can verify a ROA by bootstrapping all RIR trust anchors and recursively following prefix delegation certificates to a ROA's EE certificate. The lifetime protections (area (II)) of RPKI are enforced by definitive lifetimes of ROAs that are inherited from their certificate delegation chains (i.e., certificate lifetimes). Similar, but distinct, is that ROAs can indicate their desired cache control timings. This is, however, distinct from the object protections. Studies on lifetimes of ROAs (i.e., object security studies) are absent from the literature. Similarly, research area (III)alignment of key and object lifecycles-is absent from the protocol and the literature. While the EE certificates contained



Fig. 3. Authenticode protects objects through an X.509 delegation hierarchy.

in ROAs could (potentially) be aligned with the intended ROA lifetime, the overall verifiable lifetime of the ROA depends on the *full delegation chain* (which is decoupled from ROA lifetimes). The RPKI implements integrity protections over its objects, but does not provide a provenance mechanism (area (IV)). This, likely, is not an omission but a derivative of the fact that ROAs should not be derived from a source other than the rightful holder of a resource. Similarly, the RPKI does not have a motivation for the confidentiality of its objects (area (V)). Finally however, as with DNSSEC, the security of the *objects themselves* remains un-evaluated (area (VI)). Calculating ROV status, though important, is not equivalent to evaluating if the objects (ROAs) themselves are being protected by the RPKI system (ROV \neq ROA protections). ROV evaluates the correctness of the routing announcements, but it is not designed evaluate if the object itself is, or has remained, secure. For example, it does not have semantics to protect against newer objects that supersede older objects by the same "name," or against objects that have been otherwise invalidated. As with DNSSEC, this is not a critique of the research and does not reflect a limitation of it.

Authenticode: This protocol is designed to prove the authenticity of code objects, *i.e.*, Portable Executable (PE) files. This includes executables (.exe), dynamically loaded libraries (.dll), and drivers (.sys). It is primarily designed to verify the origin and integrity of software binaries, as depicted in Figure 3.

Authenticode attempts to implement origin authentication (area (I)) via an embedded digital signature, which ties each PE object to its "owner's certificate." However, without a formal name lookup for the PE file itself or verification of the certificate's name, it does not fully validate the origin's identity. It only verifies the binding between a PE object and the cryptographic key in the certificate. Another key aspect of Authenticode is its signature lifecycle, which connects to research area (II) (lifetime protections). Authenticode optionally supports timestamping, extending signatures' validity beyond the expiration of the signing certificates. When implemented, timestamping grants the signing certificate and signature indefinite validity. This approach aligns the lifecycle of the digital signature with the signing certificate. While timestamping addresses the issue of expiring certificates and signatures, it introduces a potential vulnerability: indefinitely

valid certificates may remain exploitable if compromised. Authenticode implements a technique to align object lifetimes with key lifecycles (area (III)) by explicitly decoupling them, by default. When a PE is created, software can obtain a signed attestation from a Time Stamp Authority (TSA) that attests that all cryptographic material was valid at creation time. Then, all certificates in the delegation chain are bundled into the PE (with their definitive expiration times included). This approach allows verifying software to use expired certificates iff they were valid at creation time, thereby "securely" decoupling their lifecycles from object lifetimes. Authenticode emphasizes integrity but makes no claims regarding provenance (research area (IV)). it ensures integrity by hashing the signed PE file. Any modification to the file after signing can be detected by comparing the computed hash during verification. Authenticode does not address the confidentiality of PE files (area (V)), focusing exclusively on owner authenticity and integrity. Similarly, the security of the PE files themselves (area (VI)) is unaddressed by the specification.

VII. DISCUSSION OF THE STATE OF THE INTERNET

In the pursuit of a generalized object security service layer, Table II illustrates that several protocol-specific incarnations of necessary security services and protections have already been successfully deployed at scale in the Internet. However, what is also deployed are sets of general-by-design object protections, lying fallow.

Object Security Services: Across the object security protocols we surveyed, roughly half operated on objects that were protocol-specific (*e.g.*, RPKI, DNSSEC, and Authenticode). However, S/MIME, PGP, and C2PA are designed for protections over generic object types. Similarly, DNSSEC implements a general namespace (the DNS') for the Internet's entire array of online services.

For origin authentication, we argue that *successful* implementations require consideration of *how* clients will be able to inter-administratively learn keys by names. In the case of RPKI, applying its namespace (*i.e.*, IP/ASN) to general objects of arbitrary types would be challenging. However, DNSSEC is built on the Internet's de facto namespace (the DNS). In fact, as outlined in our prior work [26], the DNS-based Authentication of Named Entities (DANE) [23], [24] protocol suite is designed to generalize origin authentication to other protocols. Examples include extending DANE protections for S/MIME [25] and PGP [88]. With DNSSEC's general namespace and those protocols' application to arbitrary object types, the ingredients for generalizable object security are (arguablly) already operational!

Similarly, DNSSEC, RPKI, and Authenticode include lifetime protections. Though, none of these include guidance or intuition on how to *meaningfully* set these lifetimes. DNSSEC and RPKI are soft-state protocols, which described as protecting against replay attacks. Authenticode implements this protection as optional, and off-by-default.

Interestingly, none of these protocols fully specifies the keys-to-objects protections, *except* Authenticode. In that protocol, the management of keys' lifecycles is (by default)

decoupled from objects by using Time Stamp Authorities (TSAs). The generalizability and robustness of this approach may be open to debate, but it is a rare example where this protection has been inter-administratively operationalized.

Integrity protections have been implemented by all of the protocols we considered. However, provenance has only been protected by C2PA. The type of provenance that C2PA protects involves the modification history of media objects. As we noted throughout, our view is that a more general form (process provenance) will be necessary, in order to expose the derivation history and relationships of objects. Further, as data objects now have their own *agency*, additional provenance showing objects' chain of custody (*i.e.*, where have objects been circulated) will be very important for critical thinking frameworks. The technical mechanism by which C2PA implements its provenance lends itself very well to both the protections it is designed for *and* also potential extensibility to the type of process provenance we are suggesting.

Finally, the universal lack of evaluability to the object-level protections of these protocols is a critical absence. This is, perhaps, also underscored by the expansive evaluations of the properties that are *derived* on top of these protocols (as described in Section VI). We argue that what is needed for a generalized object security service layer is an evaluation approach that can quantify *the protections on the objects themselves*.

Ways Forward: Robust object security protections will face usability challenges. The complex interplay between protections may face adoption challenges if, for users to benefit from them, they require inherit complexity. For example, if users must manually inspect process provenance trees to determine if a news story is derived from sources they trust, the protections may become too cumbersome. Moreover, if a single (secure) news story reports a view that differs from a corpus of nonsecured news stories, how should users react? For the former concern, we envision an important vein of usable security research, whereby complexity can be automated (while still being available for inspection). For example, consider a process provenance graph which shows the sources of information in its digital objects, and details where it has been promoted and seen online (including all of the embedded images and source images). Large Language Models (LLMs) have become a growing staple of Internet users who seek summaries of complex information. Providing an object secured input would allow complex interpretation of both objective and, potentially, subjective aspects of critical thinking frameworks. Furthermore, a cryptographically verifiable credential from the specific LLM used would then need to be added to a digital object's process provenance. This would still allow a user to visually inspect an object's security properties, but serve to streamline the usage. Additionally, consider training corpuses for models that online contain object secured data. In the case of the latter concern, we propose that discovering a path towards a consensus view of reality from digital objects will best be found through approaches.

A Tactical Road Map: We propose that existing generalby-design protections lay a tactical path forward. Whereas our Table II shows that no single object security protocol deployed today implements all of the security services described, several implement generalizable protections and can (and should) be combined. By using S/MIME (or PGP), DNSSEC, and C2PA together, the combined protections would include general: origin authentication, integrity, provenance, and confidentiality. Providing the linkage for this would seem to be the basic challenge, but this is precisely what the DANE protocol suite is designed for. RFCs [25], [88] specify the protocols for binding DNSSEC's origin authentication to S/MIME and PGP, respectively. These protocols also serve as ready templates for extending origin authentication protections from DNSSEC to C2PA, thereby enhancing its provenance protections. Lastly, C2PA's protocol bases its provenance protections on the general-purpose JUMBF, whose extensiblility will allow enhanced process provenance protections.

VIII. A CALL TO ARMS

We are issuing a *call to arms:* the era of unsecured digital objects on the Internet must be constrained to rapidly restore our ability to integrate *genuine* online facts to support resilient cognition and robust decision-making. To turn the tide against this Cyber-Psychosis, a first step is to develop a foundational understanding of *how* to effectively implement generalized object security in a way that supports rational critical thinking even with the prevalence of deceptive online content. Only by establishing a foundation of secure objects can we then hope to develop a cognitive sense of consensus reality from online sources and overcome Cyber-Psychosis.

We propose that the Internet needs a generalized object security service layer, to secure digital objects by default. As a straw man, in Section V we proposed six candidate types of protections needed: (I) origin authentication, (II) lifetime protections, (III) key/object lifecycle management, (IV) integrity and provenance, (V) confidentiality, and (VI) evaluability. However, we already have been deploying many of the basic pieces that we need (in some cases for decades). What stands as our challenge to the community is to use what we have, and develop any remaining pieces. Origin authentication already exists in large-scale inter-administrative protocols like DNSSEC and RPKI. What's more, DNSSEC is a generalized namespace and multiple DANE protocols have bridged it to general objects, specifically using S/MIME and PGP, as previously observed [26]. Yet, C2PA (which has produced a foundationally advanced provenance framework) has not yet embraced full origin authentication. We believe that protocols like C2PA can easily be evolved to use DANE, thereby enhancing them one step further toward full object security protocols!

Furthermore, while lifetime protections have been developed in multiple protocols, we need to backstop them with guidance about *how to use them effectively*. Our protocols that make use of public key cryptography have struggled with key-to-object lifecycle alignment. That said, Authenticode has produced one approach that manages this, and recent research in DNSSEC [63] illustrates an alternative methodology to managing this. Perhaps the final missing piece is *evaluability* of object protections. However, we contend that this is *not* a blocking precondition. The necessary basic research will be aided by pursuing holistic deployments and interdisciplinary efforts which tie an understanding into resilient critical thinking with object security and technical developments. Individually we have many of the necessary pieces, we need to reach across our disciplinary research silos and work *together!*

ACKNOWLEDGMENT

The views expressed in this work are those of the authors in their unofficial capacity and do not necessarily reflect the official policy or position of the U.S. Military Academy, U.S. Army, Office of Naval Research, or U.S. Government.

REFERENCES

- S. Dubey, "Towards finding association between decision fatigue and critical thinking in the constructive e-learning environment," *International Journal of Scientific Research and Review*, vol. 7, no. 6, pp. 2019–58, 2019.
- [2] D. Wilson, "Rop instructors' perceptions of california education code section 44910," Master's thesis, California State University, San Bernardino, 1994.
- [3] J. P. Wihbey, "The challenges of democratizing news and information: examining data on social media, viral patterns and digital influence," *Shorenstein Center Discussion Paper Series*, 2014.
- [4] D. A. Broniatowski, P. Hosseini, E. V. Porter, and T. J. Wood, "The role of mental representation in sharing misinformation online." *Journal of Experimental Psychology: Applied*, 2024.
- [5] B. Kitchens, S. L. Johnson, and P. Gray, "Understanding echo chambers and filter bubbles: The impact of social media on diversification and partisan shifts in news consumption." *MIS quarterly*, vol. 44, no. 4, 2020.
- [6] F. Zimmer, K. Scheibe, M. Stock, and W. G. Stock, "Echo chambers and filter bubbles of fake news in social media. man-made or produced by algorithms," in 8th annual arts, humanities, social sciences & education conference, 2019, pp. 1–22.
- [7] R. Thomson and C. Lebiere, "Comparing similarity and homophilybased cognitive models of influence and conformity," in *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation.* Springer, 2024, pp. 47–57.
- [8] E. Nechushtai and S. C. Lewis, "What kind of news gatekeepers do we want machines to be? filter bubbles, fragmentation, and the normative dimensions of algorithmic recommendations," *Computers in human behavior*, vol. 90, pp. 298–307, 2019.
- [9] B. Sparrow, J. Liu, and D. M. Wegner, "Google effects on memory: Cognitive consequences of having information at our fingertips," *science*, vol. 333, no. 6043, pp. 776–778, 2011.
- [10] H.-P. H. Lee, A. Sarkar, L. Tankelevitch, I. Drosos, S. Rintel, R. Banks, and N. Wilson, "The impact of generative ai on critical thinking: Selfreported reductions in cognitive effort and confidence effects from a survey of knowledge workers," in *Proceedings of the Conference on Human Factors in Computing Systems (CHI) 2025*, 2025.
- [11] R. Thomson and W. Frangia, "Investigating the use of belief-bias to measure acceptance of false information," *Computational and Mathematical Organization Theory*, pp. 1–19, 2024.
- [12] R. Thomson, C. Lebiere, J. R. Anderson, and J. Staszewski, "A general instance-based learning framework for studying intuitive decisionmaking in a cognitive architecture," *Journal of Applied Research in Memory and Cognition*, vol. 4, no. 3, pp. 180–190, 2015.
- [13] A. C. Bull, M. MacMillan, and A. Head, "Dismantling the evaluation framework," *the Library with the Lead Pipe*, pp. 551–52, 2021.
- [14] Y. Wang, Z. Su, N. Zhang, R. Xing, D. Liu, T. H. Luan, and X. Shen, "A survey on metaverse: Fundamentals, security, and privacy," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 319–352, 2022.
- [15] J. Schaad, B. Ramsdell, and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 4.0 Message Specification," IETF, RFC 8551, Apr. 2019. [Online]. Available: http://tools.ietf.org/rfc/ rfc8551.txt

- [16] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," IETF, RFC 6480, Feb. 2012. [Online]. Available: http://tools.ietf.org/rfc/rfc6480.txt
- [17] Coalition for Content Provenance and Authenticity (C2PA). Content Credentials : C2PA Technical Specification. C2pa.org: Accessed November 22, 2024. [Online]. Available: https://c2pa.org/specifications/ specifications/2.1/specs/C2PA_Specification.html
- [18] P. Mockapetris and K. J. Dunlap, "Development of the domain name system," in *Proceedings of ACM SIGCOMM*. New York, NY, USA: ACM, 1988, pp. 123–133.
- [19] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," IETF, RFC 4033, Mar. 2005. [Online]. Available: http://tools.ietf.org/rfc/rfc4033.txt
- [20] —, "Resource Records for the DNS Security Extensions," IETF, RFC 4034, Mar. 2005. [Online]. Available: http://tools.ietf.org/rfc/rfc4034.txt
- [21] —, "Protocol Modifications for the DNS Security Extensions," IETF, RFC 4035, Mar. 2005. [Online]. Available: http://tools.ietf.org/rfc/ rfc4035.txt
- [22] H. Yang, E. Osterweil, D. Massey, S. Lu, and L. Zhang, "Deploying cryptography in internet-scale systems: A case study on dnssec," *IEEE Transactions on Dependable and Secure Computing*, vol. 8, no. 5, pp. 656–669, 2010.
- [23] P. Hoffman and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA," IETF, RFC 6698, Aug. 2012. [Online]. Available: http: //tools.ietf.org/rfc/rfc6698.txt
- [24] V. Dukhovni and W. Hardaker, "The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance," IETF, RFC 7671, Oct. 2015. [Online]. Available: http: //tools.ietf.org/rfc/rfc7671.txt
- [25] P. Hoffman and J. Schlyter, "Using Secure DNS to Associate Certificates with Domain Names for S/MIME," IETF, RFC 8162, May 2017. [Online]. Available: http://tools.ietf.org/rfc/rfc8162.txt
- [26] E. Osterweil, "A cybersecurity terminarch: Use it before we lose it," *IEEE Security & Privacy*, vol. 18, no. 4, pp. 67–70, 2020.
- [27] J. Ventura, A. D. Thames, R. C. Wood, L. H. Guzik, and G. S. Hellemann, "Disorganization and reality distortion in schizophrenia: a meta-analysis of the relationship between positive symptoms and neurocognitive deficits," *Schizophrenia research*, vol. 121, no. 1-3, pp. 1–14, 2010.
- [28] W. Gibson, "Neuromancer (1984)," in *Crime and Media*. Routledge, 2019, pp. 86–94.
- [29] M. J. Steele, "Chippin'in: An analysis of the criminological concepts within cyberpunk 2077," *Games and Culture*, vol. 19, no. 2, pp. 199– 217, 2024.
- [30] A. Ross, Strange weather: Culture, science and technology in the age of limits. Verso, 1991.
- [31] B. Nyhan, "Why the backfire effect does not explain the durability of political misperceptions," *Proceedings of the National Academy of Sciences*, vol. 118, no. 15, p. e1912440117, 2021.
- [32] T. Wood and E. Porter, "The elusive backfire effect: Mass attitudes' steadfast factual adherence," *Political Behavior*, vol. 41, pp. 135–163, 2019.
- [33] B. Swire-Thompson, J. DeGutis, and D. Lazer, "Searching for the backfire effect: Measurement and design considerations," *Journal of applied research in memory and cognition*, vol. 9, no. 3, pp. 286–299, 2020.
- [34] H. M. Johnson and C. M. Seifert, "Sources of the continued influence effect: When misinformation in memory affects later inferences." *Journal of experimental psychology: Learning, memory, and cognition*, vol. 20, no. 6, p. 1420, 1994.
- [35] S. Lewandowsky, U. K. Ecker, C. M. Seifert, N. Schwarz, and J. Cook, "Misinformation and its correction: Continued influence and successful debiasing," *Psychological science in the public interest*, vol. 13, no. 3, pp. 106–131, 2012.
- [36] U. K. Ecker and L. M. Antonio, "Can you believe it? an investigation into the impact of retraction source credibility on the continued influence effect," *Memory & Cognition*, vol. 49, pp. 631–644, 2021.
- [37] M. Neijzen, "The epistemic value of doombehaviour: beyond the prudential consequences of doomscrolling, doomchecking, and doomsurfing," *Synthese*, vol. 203, no. 6, pp. 1–25, 2024.
- [38] B. Sharma, S. S. Lee, and B. K. Johnson, "The dark at the end of the tunnel: Doomscrolling on social media newsfeeds." *Technology, Mind, and Behavior*, 2022.
- [39] M. J. Crockett, "Moral outrage in the digital age," *Nature human behaviour*, vol. 1, no. 11, pp. 769–771, 2017.

- [40] S. ROY, "Mental health and cyberspace: Exploring the psychosociological needs for the new millennium," *The Posthuman Imagination: Literature at the Edge of the Human*, p. 77, 2021.
- [41] E. Stip, A. Thibault, A. Beauchamp-Chatel, and S. Kisely, "Internet addiction, hikikomori syndrome, and the prodromal phase of psychosis," *Frontiers in psychiatry*, vol. 7, p. 6, 2016.
- [42] G. S. Sales-Filho, I. D. Bandeira, N. Argollo, and R. Lucena, "Hikikomori syndrome and digital technologies: a systematic review," *Harvard review of psychiatry*, vol. 31, no. 2, pp. 50–59, 2023.
- [43] D. H. Han, Y. S. Lee, K. C. Yang, E. Y. Kim, I. K. Lyoo, and P. F. Renshaw, "Dopamine genes and reward dependence in adolescents with excessive internet video game play," *Journal of addiction medicine*, vol. 1, no. 3, pp. 133–138, 2007.
- [44] A. C. W. Huang, "Autonomic nervous system and brain circuitry for internet addiction," *Internet addiction: Neuroscientific approaches and therapeutical implications including smartphone addiction*, pp. 161–180, 2017.
- [45] J. B. Schmitt, C. A. Debbelt, and F. M. Schneider, "Too much information? predictors of information overload in the context of online news exposure," *Information, Communication & Society*, vol. 21, no. 8, pp. 1151–1167, 2018.
- [46] A. Simchon, W. J. Brady, and J. J. Van Bavel, "Troll and divide: the language of online polarization," *PNAS nexus*, vol. 1, no. 1, p. pgac019, 2022.
- [47] D. Shin, "How do people judge the credibility of algorithmic sources?" Ai & Society, pp. 1–16, 2022.
- [48] J. R. Schoenherr and R. Thomson, "When ai fails, who do we blame? attributing responsibility in human-ai interactions," *IEEE Transactions* on *Technology and Society*, 2024.
- [49] R. Thomson and J. R. Schoenherr, "Knowledge-to-information translation training (kitt): An adaptive approach to explainable artificial intelligence," in *Adaptive Instructional Systems*, R. A. Sottilare and J. Schwarz, Eds. Cham: Springer International Publishing, 2020, pp. 187–204.
- [50] A. Kozyreva, S. Lewandowsky, and R. Hertwig, "Citizens versus the internet: Confronting digital challenges with cognitive tools," *Psychological Science in the Public Interest*, vol. 21, no. 3, pp. 103–156, 2020.
- [51] G. Liu, "Moving up the ladder of source assessment: Expanding the craap test with critical thinking and metacognition," *College & Research Libraries News*, vol. 82, no. 2, p. 75, 2021.
- [52] J. R. Schoenherr and R. Thomson, "Persuasive features of scientific explanations: Explanatory schemata of physical and psychosocial phenomena," *Frontiers in Psychology*, vol. 12, 2021.
- [53] M. Soprano, K. Roitero, D. La Barbera, D. Ceolin, D. Spina, G. Demartini, and S. Mizzaro, "Cognitive biases in fact-checking and their countermeasures: A review," *Information Processing & Management*, vol. 61, no. 3, p. 103672, 2024.
- [54] I. Ajzen, "The theory of planned behavior," Organizational behavior and human decision processes, vol. 50, no. 2, pp. 179–211, 1991.
- [55] D. Charnock, S. Shepperd, G. Needham, and R. Gann, "Discern: an instrument for judging the quality of written consumer health information on treatment choices." *Journal of Epidemiology & Community Health*, vol. 53, no. 2, pp. 105–111, 1999.
- [56] C. Boyer, M. Selby, J.-R. Scherrer, and R. D. Appel, "The health on the net code of conduct for medical and health websites," *Computers in biology and medicine*, vol. 28, no. 5, pp. 603–610, 1998.
- [57] M. Caulfield, "Information literacy for mortals. pil provocation series. volume 1, number 5." *Project Information Literacy*, 2021.
- [58] A. Tardiff, "Have a ccow: A craap alternative for the internet age," *Journal of Information Literacy*, vol. 16, no. 1, 2022.
- [59] D. Sye and D. S. Thompson, "Tools, tests, and checklists: The evolution and future of source evaluation frameworks," *J. New Librarianship*, vol. 8, p. 76, 2023.
- [60] I. A. Portillo, C. V. Johnson, and S. Y. Johnson, "Quality evaluation of consumer health information websites found on google using discern, craap, and honcode," *Medical reference services quarterly*, vol. 40, no. 4, pp. 396–407, 2021.
- [61] E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.3," IETF, RFC 8446, Aug. 2018. [Online]. Available: http: //tools.ietf.org/rfc/rfc8446.txt
- [62] L. J. Osterweil, L. A. Clarke, A. M. Ellison, R. Podorozhny, A. Wise, E. Boose, and J. Hadley, "Experience in using a process language to define scientific workflow and generate dataset provenance," in *Proceedings of the 16th ACM SIGSOFT International Symposium* on Foundations of Software Engineering, ser. SIGSOFT '08/FSE-16. New York, NY, USA: Association for Computing Machinery, 2008, p. 319–329. [Online]. Available: https://doi.org/10.1145/1453101.145314

- [63] E. Osterweil, P. F. Tehrani, T. C. Schmidt, and M. Wählisch, "From the Beginning: Key Transitions in the First 15 Years of DNSSEC," *Transactions on Network and Service Management (TNSM)*, vol. 19, no. 4, pp. 5265–5283, December 2022.
- [64] Y. L. Simmhan, B. Plale, and D. Gannon, "A survey of data provenance in e-science," ACM Sigmod Record, vol. 34, no. 3, pp. 31–36, 2005.
- [65] S. S. Sahoo and A. P. Sheth, "Provenir Ontology: Towards a Framework for eScience Provenance Management," The Ohio Center of Excellence in Knowledge-Enabled Computing, Kno.e.sis Publication 80, 2009. [Online]. Available: https://corescholar.libraries.wright.edu/knoesis/80/
- [66] S. Miles, P. Groth, M. Branco, and L. Moreau, "The requirements of using provenance in e-science experiments," *Journal of Grid Computing*, vol. 5, pp. 1–25, 2007.
- [67] J. Zhao, C. Wroe, C. Goble, R. Stevens, D. Quan, and M. Greenwood, "Using semantic web technologies for representing e-science provenance," in *International Semantic Web Conference*. Springer, 2004, pp. 92–106.
- [68] L. Rosenthol, "C2pa: the world's first industry standard for content provenance (conference presentation)," in *Applications of Digital Image Processing XLV*, vol. 12226. SPIE, 2022, p. 122260P.
- [69] N. Fotos and J. Delgado, "Ensuring privacy in provenance information for images," in 2023 24th International Conference on Digital Signal Processing (DSP). IEEE, 2023, pp. 1–5.
- [70] J. Sedlmeir, A. Rieger, T. Roth, and G. Fridgen, "Battling disinformation with cryptography," *Nature Machine Intelligence*, vol. 5, no. 10, pp. 1056–1057, 2023.
- [71] S. Petrangeli, H. Wang, M. Fisher, D. Kozma, M. Mahamli, P. Blumenthal, and A. Parsons, "Integrating content authenticity with dash video streaming," in *Proceedings of the 15th ACM Multimedia Systems Conference*, ser. MMSys '24. New York, NY, USA: ACM, 2024, p. 492–498. [Online]. Available: https: //doi.org/10.1145/3625468.3652198
- [72] P. Zimmerman, "The official pgp user's guide. prz@ acm. org, 1994."
- [73] Microsoft. (2008, March) Windows Authenticode Portable Executable Signature Format. Statistica: Accessed October 20, 2024.
 [Online]. Available: http://download.microsoft.com/download/9/c/5/ 9c5b2167-8017-4bae-9fde-d599bac8184a/Authenticode_PE.docx
- [74] W3C. (2024, October) Verifiable Credentials Data Model v2.0. W3.org: Accessed November 22, 2024. [Online]. Available: https: //www.w3.org/TR/vc-data-model-2.0/
- [75] E. Osterweil. Secspider. [Online]. Available: https://secspider.net/
- [76] (2024, November) Use of DNSSEC Validation for World (XA). Stats.labs.apnic.net: Accessed November 12, 2024. [Online]. Available: https://stats.labs.apnic.net/dnssec/XA
- [77] Y. Kocaogullar, E. Osterweil, and L. Zhang, "Towards a decentralized internet namespace," in *Proceedings of the ACM Conext-2024 Workshop* on the Decentralization of the Internet, ser. DIN '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 36–41. [Online]. Available: https://doi.org/10.1145/3694809.3700746
- [78] M. Müller, M. Thomas, D. Wessels, W. Hardaker, T. Chung, W. Toorop, and R. v. Rijswijk-Deij, "Roll, Roll, Roll Your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover," in *Proc. of ACM IMC '19*. New York, NY, USA: ACM, 2019, pp. 1–14.
- [79] Q. A. Chen, E. Osterweil, M. Thomas, and Z. M. Mao, "Mitm attack by name collision: Cause analysis and vulnerability assessment in the new gtld era," in 2016 IEEE Symposium on Security and Privacy (SP). IEEE, 2016, pp. 675–690.
- [80] Q. A. Chen, M. Thomas, E. Osterweil, Y. Cao, J. You, and Z. M. Mao, "Client-side name collision vulnerability in the new gtld era: A systematic study," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 941–956.
- [81] T. Chung, R. van Rijswijk-Deij, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "A longitudinal, end-to-end view of the DNSSEC ecosystem," in *Proc. of 26th USENIX Security Symposium*. Washington D.C., USA: USENIX Assoc., 2017, pp. 1307–1322.
- [82] W. Lian, E. Rescorla, H. Shacham, and S. Savage, "Measuring the practical impact of DNSSEC deployment," in *Proc. of the 22nd USENIX Security Symposium.* Washington D.C., USA: USENIX Assoc., 2013, pp. 573–588.
- [83] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," IETF, RFC 6811, Jan. 2013. [Online]. Available: http://tools.ietf.org/rfc/rfc6811.txt
- [84] N. Rodday, Í. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, "The Resource Public Key Infrastructure (RPKI): A Survey on Measurements and Future Prospects," *IEEE*

Transactions on Network and Service Management, vol. 21, no. 2, pp. 2353–2373, April 2024.

- [85] W. Li, Z. Lin, M. I. Ashiq, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "RoVista: Measuring and Analyzing the Route Origin Validation (ROV) in RPKI," in *Proceedings of the 2023* ACM Internet Measurement Conference, ser. IMC '23. New York, NY, USA: ACM, 2023, p. 73–88. [Online]. Available: https://doi.org/10.1145/3618257.3624806
- [86] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "Ripki: The tragic story of rpki deployment in the web ecosystem," in *Proceedings of the 14th ACM Workshop on Hot Topics* in Networks, 2015, pp. 1–7.
- [87] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, "Towards a rigorous methodology for measuring adoption of rpki route validation and filtering," *SIGCOMM Comput. Commun. Rev.*, vol. 48, no. 1, p. 19–27, Apr. 2018. [Online]. Available: https://doi.org/10.1145/3211852.3211856
- [88] P. Wouters, "DNS-Based Authentication of Named Entities (DANE) Bindings for OpenPGP," IETF, RFC 7929, Aug. 2016. [Online]. Available: http://tools.ietf.org/rfc/rfc7929.txt

Robert Thomson is an Associate Professor in the Engineering Psychology Program at the United States Military Academy, a Principal Research Scientist at the Army Cyber Institute, and currently serves as the Dean's Fellow for Research. Dr. Thomson has over 13 years of post-graduate research experience and over 75 invited and refereed academic publications in the domains of computational modeling, intelligence analysis, cybersecurity, and artificial intelligence. He has been selected and supported more than \$25M in reimbursable research from IARPA, DARPA, NIH, the Office of Naval Research, and several Army organizations.

Quan Nguyen is a second-year Ph.D. student in Computer Science at the George Mason University (GMU) under the advising of Prof. Eric Osterweil (GMU). He is interested in topics in Cryptography and Internet Security. His research focuses on evaluating the longitudinal security of current Internet's object security protocols, e.g. DNSSEC, etc., and analyzing the lessons from the deployments of those protocols that can guide the design and implementation of the future Internet's object security protocols.

Essien Ayanam is a Ph.D. student in Computer Science at George Mason University (GMU). He has over 15 years of experience in designing and deploying secure network architectures, security policies, and procedures. He serves as a Senior Network Security Engineer in the Investigative Technology Division at the Drug Enforcement Administration, addressing critical security challenges in the government space. His research focus is on measurable security for Internet-based protocols, digital object protections, and automotive network technology (CAN). As a researcher at GMU's Measurable Security Lab, he examined the impact of Internet protocols on security mechanisms and assisted with the development of a framework for evaluating digital object protections.

Matthew Canham is a former Supervisory Special Agent with the Federal Bureau of Investigation (FBI), he has a combined twenty-one years of experience in conducting human-technology and security research. He currently holds an affiliated faculty appointment with George Mason University, where his research focuses on threats posed by maliciously produced AI generated content and synthetic media social engineering. Dr. Canham recently founded the Cognitive Security Institute, a non-profit organization dedicated to understanding the key components of cognitive attacks and discovering the best ways to defend against these.

Thomas C. Schmidt is a Professor of computer networks and Internet technologies with the Hamburg University of Applied Sciences, Hamburg, Germany, where he heads the Internet Technologies Research Group. He was the Principal Investigator in a number of EU, nationally funded, and industrial projects, as well as a Visiting Professor with the University of Reading, Reading, U.K. He is a cofounder and a coordinator of several successful open source communities. His current research interests include development, measurement, and analysis of large-scale distributed systems, such as the Internet or its offsprings.

Matthias Wählisch is a Full Professor and holds the Chair of Distributed and Networked Systems at the Faculty of Computer Science at TU Dresden. He is also a Research Fellow of the Barkhausen Institut. His research and teaching focus on scalable, reliable, and secure Internet communication. This includes the design and evaluation of networking protocols and architectures, as well as Internet measurements and analysis. Matthias is involved in the IETF since 2005 and co-founded multiple successful open source projects such as RIOT and the RTRlib. Eric Osterweil is the Associate Director of the Center for Assurance Research and Engineering (CARE) and an Assistant Professor in the Department of Computer Science at George Mason University. He is the former co-Chair of the 2nd Security, Stability, and Resiliency (SSR2) Review Team in the ICANN Community. His research focuses on Internet-scale object security, cybersecurity protocols, and large-scale Internet measurements. He has authored and co-authored over 130 papers, articles, patents, and RFCs and maintains the most complete longitudinal measurement corpus of the global deployment of DNSSEC, whose data spans from the protocol's final standardization in 2005 through to today.