DISSERTATION

AN ALGORITHMIC SEMANTIC ANALYSIS OF CYBER SECURITY AND RESILIENCE GUIDANCE AGAINST INTERDISCIPLINARY UNDERSTANDING OF RESILIENCE CONCEPTS ACROSS TIME AND SCALE

Submitted by

Ryan Hilger

Department of Systems Engineering

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2025

Doctoral Committee:

Advisor: Steve Simske

Jennifer Cross
Jeremy Daily
Indrakshi Ray

ABSTRACT

AN ALGORITHMIC SEMANTIC ANALYSIS OF CYBER SECURITY AND RESILIENCE GUIDANCE AGAINST INTERDISCIPLINARY UNDERSTANDING OF RESILIENCE CONCEPTS ACROSS TIME AND SCALE

This dissertation bridges critical gaps between cybersecurity frameworks and interdisciplinary resilience theory through innovative algorithmic analysis. Rather than pursuing an elusive singular definition of resilience, I employ statistical modeling and machine learning techniques to extract core resilience attributes from a diverse corpus of 102 unique definitions across fields including ecology, psychology, disaster management, and organizational studies. My research addresses two fundamental questions: (1) Does any existing cybersecurity strategy or guidance document comprehensively address resilience across temporal and scalar dimensions? (2) How do current frameworks conceptualize and operationalize resilience?

The methodological approach integrates term frequency-inverse document frequency (tf*idf), Latent Dirichlet Allocation, and bidirectional encoder representations from transformers (BERT) algorithms to construct a novel classification scaffold based on time and scale dimensions. This scaffold systematically evaluates 37 cybersecurity frameworks and 12 non-cyber resilience frameworks against core resilience attributes. Results reveal significant gaps between cybersecurity guidance and interdisciplinary resilience concepts, with most frameworks focusing predominantly on technical and sociotechnical aspects while neglecting broader organizational, community, and temporal dimensions of resilience.

This research makes several key contributions: (1) establishing a data-driven classification framework for assessing resilience features in guidance documents, (2) demonstrating that no single existing framework adequately addresses resilience across all relevant dimensions, and (3) providing a foundation for developing more comprehensive cyber resilience strategies. The findings offer both theoretical advancement in conceptualizing cyber resilience and practical guidance for organizations seeking to build more adaptable and resilient systems across multiple time horizons and organizational scales.

ACKNOWLEDGEMENTS

The seed of this idea was first planted in conversations with Dr. Robert Templeman in early 2020 while discussing how to make major defense weapons systems, particularly large system of systems ones, more resistant or resilience to cyberattacks from across the industrial base. It became rapidly apparent that most of the community around me did not have the conceptual framework from which to grasp how cyber vulnerabilities or attacks might propagate across the ecosystem that supports a weapon system, or how to even begin approaching it. I had voiced to Dr. Templeman my desire to eventually pursue a doctorate in a field, and he quickly told me I already had the kernel of a topic. From there the flood gates opened.

I spent several months during the application process conducting informational interviews with various people from around the Department of Defense who were involved in cyber in some fashion Dr. Evan Austin, Dr. Thomas Llanso, Dr. Thomas Rondeau, Dr. Raymond Richards, Dr. Sergey Bratus, Dr. Walter Weiss, Dr. William “Dollar” Young, Brad Martin, Dan Ragsdale, Sara Standard, and Tyson Meadors. Their early guidance on both the subject matter and on getting a doctorate were invaluable in preparing me for this journey. From an industry perspective, Alex Sharpe’s timely feedback and insights throughout the process provided a crucial perspective. Dr. Igor Linkov and Dr. Alexander Kott, two heavyweights in the field of cyber resilience, were generous with their time, feedback, and mentoring, and their body of research was a fountain of information and insights.

Part of this research was supported by the Cyber Statecraft Initiative at the Atlantic Council, under the leadership and expertise of Trey Herr. Him and his team, mainly Safa Shahfan Edwards and Stewart Scott, took a chance on a relative newcomer to the cybersecurity field and

gave me an environment to explore how ecosystems, resilience, complex adaptive systems, and other interdisciplinary theories could inform cyber resilience efforts. I am humbled and grateful for their generous support and funding for this research. Similarly, to Don Goff and Dan Wolf of Cyber Pack Ventures, supporting Brad Martin and the team at the Science of Security Virtual Organization's Cybersecurity in Contested Computing Environments Virtual Institute, for their generous funding and opportunity to share aspects of my research into trust and observability in cyber ecosystems with a community of peers.

To my committee, Drs. Jeni Cross, Jeremy Daily, and Indrakshi Ray, who challenged me at every opportunity to ask more questions, find ways to quantify what I was doing, and to put more rigor behind my analyses.

Dr. Steve Simske, my advisor and true mentor. You've been rock solid throughout this process. I could not have asked for a better advisor and mentor. You were always unbelievably prompt with feedback, document reviews, and entertaining lines of questioning from me. When we first met in person in late 2021 at Avogadro's Number in Fort Collins, I knew I had met a like mind and kindred spirit—carrying the book *Salt* by Mark Kurlansky gave it away. Thank you for taking a chance on me and for your support and mentorship over the last four years, and for the many years of partnership yet to come. I admired your story in transitioning from Hewlett Packard Labs to academia, and I seek to carry out my own version of that someday.

Lastly and most importantly, an unending stream of thanks and love to my beautiful wife, Heather. She may still ask for the thousandth time what I'm writing my dissertation on, but her love and support has been utterly unwavering. Her steadfast grace and flexibility in finding new ways to integrate my work, homeschooling our kids, our home, and our academic lives never ceases to amaze me, and made this journey of mine possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1: THE GLOBAL THREAT ENVIRONMENT AND THE CASE FOR CYBER RESILIENCE

The world feels like a decidedly more dangerous place with each passing year. Individuals, organizations, and nations face a wide array of challenges that threaten our mental well-being, our ability to create and sustain, our capacity to cope with change and stress, and even our way of life: climate change, threatened supply chains and food insecurity, increasing political and sectarian violence, the specter of a major war or three between nuclear-armed powers, the onslaught in cyberspace, and so many others. In this environment, survival depends on our individual and collective ability to bounce back, to adapt, and to find new ways to cope with our current challenges. Humanity has been doing this for millennia, though the needed pace of adaptation seems to be accelerating as technology pushes the demands on our cognitive systems toward machine speeds. Our ability to adapt and evolve has not quite kept pace, but in lieu of that, humanity has a profound capacity for adaptation and resilience that helps us cope with the accelerating changes around us.

The rise of the internet and the exponential increase in the complexity of the modern world challenges our mental models for how we think the world works and how it ought to work. The computing age and the recent explosion of artificial intelligence and machine leaning capabilities strains the human capacity for understanding and coping with changes and threats to our world. Evolution has not allowed us, yet, to cope with changes occurring at machine speed and from avenues we cannot see, hear, taste, smell, or touch—the sensory capacity that helps us make sense of our world. The cyber environment is tough to visualize, comprehend, and respond to.

Cybersecurity generally began in response to the Morris worm in 1988, the first documented case of what today we would term a cyberattack with malware, though it was intended as a harmless tool to map the growing internet. Since then, both the industry and the discipline of cybersecurity have grown tremendously in response to the increasing complexity of computer networks, the growth of the internet, and the ability of people to exploit that complexity to achieve sinister gains. The explosive growth of software in nearly every aspect of human life has given cyber actors a great surface to attack, and a more challenging and complex surface to defend. Hackers continue to do what they have always done: find and exploit vulnerabilities in software that allow them to achieve their desired aims, whether financial, political, or otherwise.

Defenders, on the other hand, must balance the need for information technology resources and software to give an organization competitive advantage with the need to keep that information and those systems protected from unauthorized access and undesired behaviors. As the complexity of software and cybersecurity continues to grow, defenders are inundated with new cybersecurity products designed to help them convert the deluge of data coming in at machine speeds into actionable information that people can do something with at their speed.

But as the complexity of the systems that cyber defenders must manage increases, their ability to fully understand the system of systems and its myriad behaviors decreases [1-5]. The complex system of systems, or the complex sociotechnical system, has evolved to the point of no longer being completely knowable by a single individual. That threshold has significant ramifications for defense and both technical and sociotechnical system functioning. As software, inclusive of cloud computing, becomes the central enabling technology that governs how we work, how we deliver value, and how we run our lives, cyber defenders must design, integrate,

provision, and protect the systems that help us generate those outcomes. The increased number of users, both internal and consumers, that interact with those systems has grown tremendously over the last few decades. Current cybersecurity guidance treats people internal to the organization, those accessing the systems to create value and do the work, as users to be trained and monitored for appropriate system use and resistance to social engineering tactics. Each customer becomes a potential threat vector for exploitation, whether through legitimate or illegitimate access. News reports of data breaches and other cyber-attacks undermine our capacity for trust in these systems.

The guidance available to the cybersecurity community provides the best practices to secure systems, plan for and respond to incidents, manage user trust and access, and myriad other topics. However, that guidance does not extend into the sociotechnical domain to help organizations understand *how* they should be organized to do the work or minimize vulnerabilities. In 1968, Melvin Conway applied this concept in his classic paper "How do committees invent?" The paper offered the often-quoted law today, known as Conway's Law: "to the extent that an organization is not completely flexible in its communication structure, that organization will stamp out an image of itself in every design in produces" [1]. While this statement may seem pithy, it reveals a far greater truth about sociotechnical design and cyber vulnerabilities. The interfaces between which components, systems, and organizations communicate often become the entry point for hackers to try to gain control of the system. This includes items not developed within the organization, such as leveraging commercial software and open source software. Industry relates this discipline to mapping the attack surface, but the methods for doing so make it difficult to capture the complexity of sociotechnical systems design, let alone keep up with it as the system evolves daily, either intentionally or

unintentionally. Adversaries in cyberspace understand this implicitly and continually probe the seams for entry points, from which they can escalate privileges or move laterally through the system and begin conducting reconnaissance or attacks on the target, and each adversary's motives are different.

Security, Robustness, and Resilience

The global cyber threat environment challenges organizations, especially those in critical infrastructure sectors, to operate without setbacks. Complicating this challenge is the cybersecurity industry itself. While well-intentioned, the cybersecurity industry today uses language in marketing and business development that muddies the water further. Often conflating terms such as security, resilience, reliability, and robustness, the misuse of those terms may lead organizations to believe they are more secure, reliable, or resilient than they are. Definitions matter, and a firm foundation is needed to understand how terms such as security, resilience, reliability, robustness, and many others, overlap and support each other. The academic literature on resilience generally fails to agree on a single definition for resilience broadly, let alone cyber resilience [2-9]. To be fair, the terms are related, but they are not synonymous.

Security, robustness, reliability, and resilience work in conjunction with each other to describe features and phases in a complex adaptive system. From a cyber perspective, security involves the actions necessary to prevent a cyberattack in the first place. Robustness is the ability of the system to continue operating at prescribed levels while under attack. Reliability is another facet of robustness in ensuring that the system can degrade gracefully or continue operating for the required durations without failure. Resilience is generally the ability to adapt in the aftermath

of an attack. Resilience, as will be described in Chapter 3, is thus an emergent property of a complex adaptive system. The technical systems in a cyberattack—those operating without the human component—will not develop or exhibit emergent behaviors in response. For example, a firewall can only block signals as defined by its rule set. But a human operating a firewall becomes capable of emergent behaviors and can respond to a successful penetration by updating the rule set and making the firewall harder to penetrate in the next attack.

The cybersecurity industry produces standards, frameworks, strategies, and other guiding documents to help organizations, communities, sectors, and other ecosystems improve their cyber readiness. These documents have begun to address resilience, but, as this research seeks to understand, it remains to be seen if the usage of resilience in the documents actually fits with what resilience is and encodes the features of resilience into the document to produce resilient outcomes for the organization using it.

Research Questions and Expected Contributions

This research seeks to determine if existing literature, frameworks, standards, or other guidance use resilience and its concepts properly as the interdisciplinary community understands it. This includes an in-depth discussion of resilience as an interdisciplinary concept and seeks to determine if the existing literature conflates resilience with other terms. Conflation of terminology would make it more challenging for individuals and organizations to understand how to create the conditions for resilience to emerge in their organization and the broader ecosystem surrounding it. The key research questions and associated hypothesis are briefly

presented below. Chapter 2 provides a detailed discussion of the research questions, hypotheses, falsification criteria, and the research design to evaluate these hypotheses.

Research Questions and Hypotheses

Research Question 1 (RQ1): Does a single strategy or guidance document exist to provide organizations at multiple levels of society with support for improving cyber resilience?

Hypothesis 1 (H1): If no single document exists that addresses most or all aspects of resilience, then organizations will not have sufficient guidance to develop complete strategies to improve cyber resilience. Put plainly, does a document exist today that satisfies RQ1? The foundational assumption for this work is that this document does not exist.

Hypothesis 2 (H2): If one or more documents exist that address some aspects of resilience, then they can be leveraged and expanded to create a single document for organizations to improve cyber resilience.

Research Question 2 (RQ2): How do existing documents address resilience?

Hypothesis 3 (H3): If the existing documents incorporate most or all aspects of resilience, then there should be significant overlap with the interdisciplinary research and concepts on resilience.

Hypothesis 4 (H4): If the existing documents focus primarily on the technical and sociotechnical aspects of cybersecurity, then the documents will lack sufficient guidance on most or all aspects of resilience.

Expected Contributions

This research will produce the following contributions. First, a classification framework from which future frameworks or guiding documents can be assessed for features of resilience—cyber or otherwise. Second, an assessment of the existing frameworks and the extent to which the framework addresses some or all aspects of resilience as the interdisciplinary community understands it. Third, it will identify the gaps in the frameworks, or combined set of frameworks, to show where the cybersecurity community can focus future efforts on developing guidance to support all aspects of resilience, thus giving organizations the ability to develop strategies to improve cyber resilience.

CHAPTER 2: RESEARCH DESIGN AND METHODOLOGY

Research Questions and Hypotheses

In the context of the global cyber and geopolitical environment, organizations at all levels, from individuals and small businesses to national governments and international institutions, have a pressing need to increase their cyber resilience. Compounding this challenge is the ambiguity around the idea of organizational resilience and application of resilience principles in a cybersecurity context [2]. This study seeks to understand if the organizations at all levels have sufficient cybersecurity guidance, standards, or other frameworks to improve cyber resilience. The following research questions and hypotheses, shown in Figure 1**Error! Reference source not found.** below, provide a structured approach to understanding if society has sufficient resources to achieve this goal; and if not, to understand where guidance can improve based on identified gaps, definitions, and integrating interdisciplinary research.

Figure 1: Structure of the research questions, hypotheses, and falsification criteria

The first research question seeks to find if any of the myriad cyber security documents, frameworks, or strategies available today encompass the attributes of resilient system as identified by interdisciplinary research. It prompts two additional questions. First, how do the existing documents address resilience? Second, how the existing guidance might be combined to form a document that would enable organizations at multiple levels to grapple with improving cyber resilience. The hypotheses and their falsification criteria follow. Given the complicated structuring of the research questions and hypotheses, stating explicit falsification criteria will better support the final evaluation of the hypotheses.

Research Question 1: Do existing documents provide scalable support for improving cyber resilience?

- RQ1: Does a single strategy or guidance document exist to provide organizations at multiple levels of society with support for improving cyber resilience?
- Hypothesis 1 (H1): If no single document exists that addresses most or all aspects of resilience, then organizations will not have sufficient guidance to develop complete strategies to improve cyber resilience. Put plainly, does a document exist today that satisfies RQ1? The foundational assumption for this work is that this document does not exist.
    - Falsification 1 (F1): A document exists that addresses most or all aspects of resilience as identified in the work to test Hypothesis 3 (H3).
- Hypothesis 2 (H2): If one or more documents exist that address some aspects of resilience, then they can be leveraged and expanded to create a single document for organizations to improve cyber resilience.
    - Falsification 2 (F2): Hypothesis 1 (H1) is proven true.

Research Question 2: How do existing documents address resilience?

- RQ2: How do existing documents address resilience?
- Hypothesis 3 (H3): If the existing documents incorporate most or all aspects of resilience, then there should be significant overlap with the interdisciplinary research and concepts on resilience.
    - Falsification 3 (F3): The analysis indicates gaps between the interdisciplinary concepts on resilience and the existing documents.

- Hypothesis 4 (H4): If the existing documents focus primarily on the technical and sociotechnical aspects of cybersecurity, then the documents will lack sufficient guidance on most or all aspects of resilience.
    - Falsification 4 (F4): The analysis indicates few to no gaps in one or more existing documents.

This set of research questions, hypotheses, and falsification criteria presents a challenge in understanding the differences between them, where they overlap, and other nuances. To help provide an explicit understanding of each hypothesis in the context of the broader body of literature on cybersecurity, resilience, sociotechnical ecosystems, and associated topics, the corpus can be presented visually, as shown in Figure 2, below.

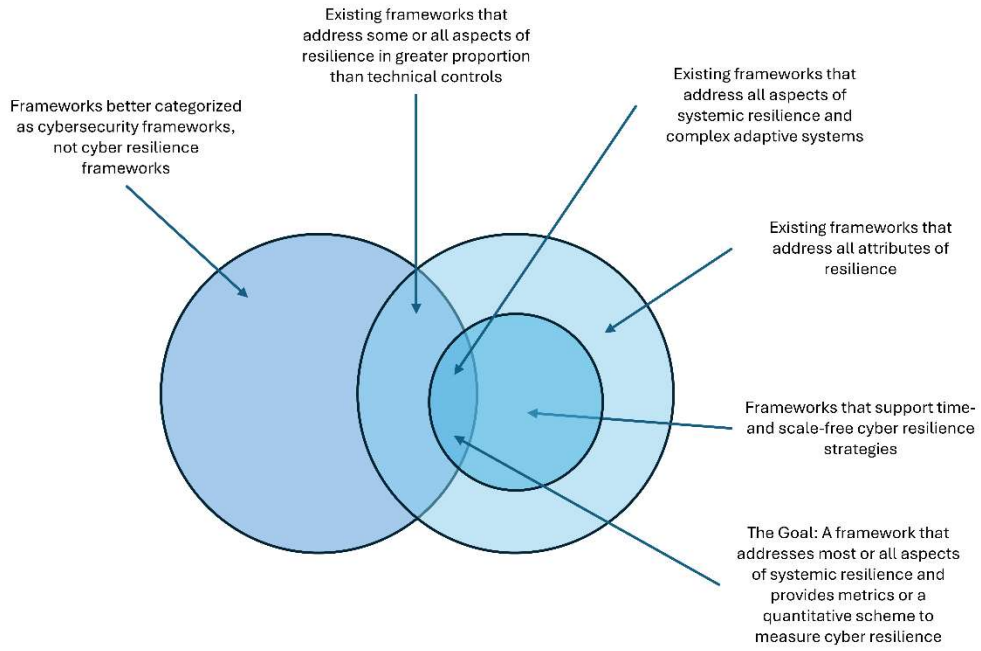

Figure 2: A Venn diagram showing the broad corpus of the research into cybersecurity and resilience frameworks, the subset of resilience frameworks that support time- and scale-free resilience strategies, and their overlap.

Research Design

The research design for this study employs an integrative, multi-phase approach that combines both qualitative and quantitative methods to address the research questions concerning the adequacy of existing cybersecurity frameworks in promoting cyber resilience [10]. The four-phase structure, as shown in Figure 3, is designed to systematically investigate the gaps and overlaps within current frameworks, leveraging a combination of interdisciplinary literature review, machine learning algorithms, and quantitative and qualitative analysis. This approach is particularly well-suited to the study's goals because it allows for a comprehensive examination of both technical and sociotechnical aspects of resilience, drawing on insights from various disciplines such as ecology, psychology, resilience engineering, complex adaptive systems, and cybersecurity.

In the first phase, the interdisciplinary literature review provides the foundation for identifying key resilience attributes, which will be codified into a matrix that serves as the scaffold for the entire analysis. Subsequent phases focus on applying statistical modeling and algorithmic techniques like term frequency-inverse document frequency (tf*idf) and bidirectional encoder representations from transformers (BERT) to classify and analyze existing cybersecurity frameworks against the matrix. These methods are chosen for their ability to process large textual datasets and identify patterns in term usage and framework features. For example, tf*idf allows for the identification of significant terms that may point to critical resilience attributes, while BERT offers deeper insights into the semantic relationships between terms across frameworks. This systematic, algorithm-driven approach ensures that both the technical rigor and conceptual breadth of resilience are captured, addressing the hypotheses related to the comprehensiveness and scalability of existing frameworks.

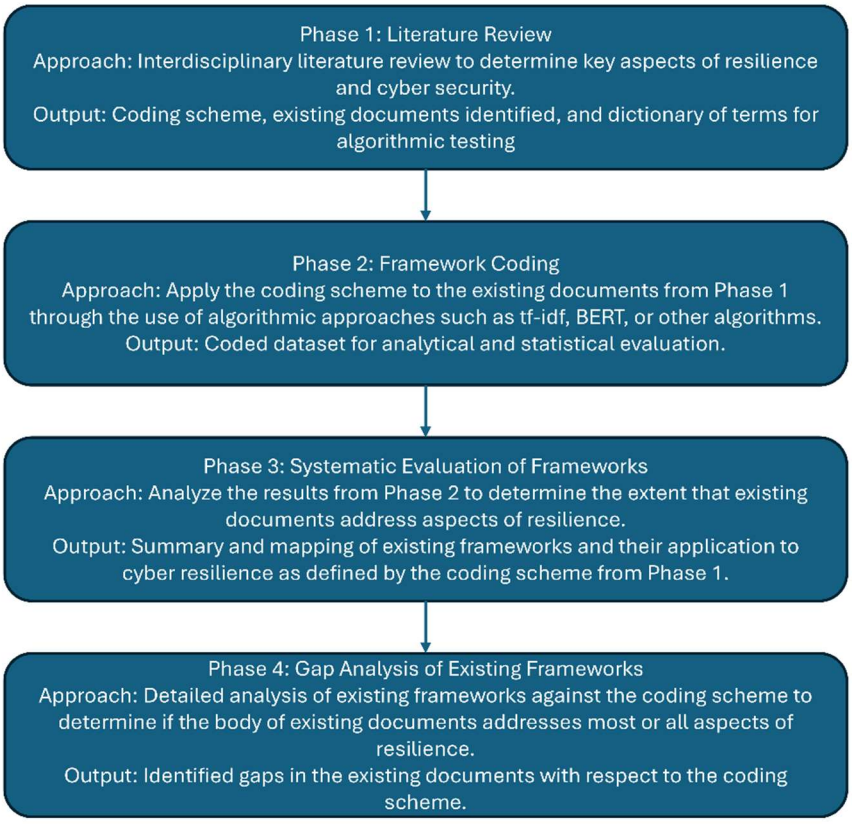

Figure 3: The four-phase research plan with hypothesis mapping to develop a broadly applicable cyber resilience framework

First, an interdisciplinary literature review to determine the key aspects of resilience and how they relate to cybersecurity will be used to build a coding scheme and dictionary of terms to support algorithmic approaches to assessing the existing body of cybersecurity or cyber resilience frameworks. Second, the existing documents will be analyzed with the coding scheme using algorithmic approaches, such as term frequency-inverse document frequency (tf*idf), BERT, or other algorithms. This phase tests hypotheses three (H3) and four (H4) and partially addresses the second research question (RQ2). The output of this phase will be a coded dataset that can be used for statistical and analytical evaluation. Using that output, the third phase will systematically evaluate the existing frameworks against the dictionary and coding scheme from

phase one to map the frameworks against the derived aspects of resilience. This phase completes the assessed work needed to address the second research question (RQ2). Phase four will conduct a detailed analysis of the results of phases one and three for gaps in the existing frameworks against some or all aspects of resilience. This phase will produce a gap analysis that directly answers the first hypothesis (H1). The remainder of this chapter provides a detailed description of the approach to data curation common to much of the analysis and the five phases of the research plan, their objectives, and their outputs.

The research design carefully integrates the hypotheses with the selected methods to ensure that each research question is systematically addressed. For instance, Hypothesis 1 (H1), which posits that no single document encompasses all aspects of resilience, is directly tested through the gap analysis in Phases Four and Five. The coding scheme developed in Phase One, based on the interdisciplinary literature review, provides a structured way to assess existing frameworks against resilience attributes, allowing for the identification of gaps that would support or refute this hypothesis. Hypothesis 2 (H2), which suggests that multiple frameworks could be combined to address these gaps, is explored through the quantitative and qualitative analysis of existing frameworks in Phase Three. By mapping key resilience features from various documents into the classification matrix, the research can determine whether combining elements from different sources could create a comprehensive resilience framework.

Hypotheses 3 and 4 (H3 and H4), which focus on how well existing frameworks address both technical and broader resilience attributes, are evaluated using machine learning models such as tf*idf and BERT in Phases Two and Three. These models allow for a detailed analysis of the language and concepts used in the frameworks, allowing for a determination of whether technical aspects are overemphasized at the expense of sociotechnical resilience principles. This

tight alignment between hypotheses and methods ensures that the research process remains focused and that each phase of analysis is directly linked to testing specific hypotheses.

Data Curation

The data needed to conduct the research proposed comes from several sources and can be divided broadly into two categories: 1) the interdisciplinary literature on resilience, which will be needed to inform the work in Phases One and Two, and 2) the body of existing cybersecurity and resilience frameworks to analyze in Phases Three and Four. This section will describe how the data for category will be collected, processed, classified, and used since the methods for both analyses are generally common.

Data Collection for Literature Review and Scaffold Development

Literature for the literature review and scaffold development was found through a systematic search using the SCOPUS database, Google Scholar, internet searches, and materials referenced by literature found from those means. Keywords for the searches included various combinations of the following words: cyber, resilience, ecosystems, networks, cybersecurity, cyberattacks, network science, resilience engineering, complex adaptive systems, adaptive capacity, graceful extensibility, high reliability organizations, safety culture, network trust, transitive trust, cyber supply chain, risk management, cyber policy, personal resilience, organizational resilience, and others. Literature includes peer reviewed journal articles, published reports from various organizations, government policies, books, and corporate websites. 144 articles, books, or other sources were identified. The search should not be considered exhaustive

since some potential articles were excluded on the basis of inaccessibility resulting from lack of a publicly available document or subscription access to the source.

Data Collection for Existing Cybersecurity and Resilience Frameworks

The work of Phases Three and Four to address RQ2 requires the curation of existing frameworks to assess against the existing interdisciplinary literature. The existing frameworks were found primarily through internet searches and informal discussions with subject matter experts in the cybersecurity field. Keywords for the searches included various combinations of cyber resilience framework, index, cyber resilience assessment, cybersecurity framework, and others. 37 documents were identified across a range of cybersecurity topics ranging from narrow technical guidance to organizational cyber risk management to threat assessment to sector specific and national cyber standards. As a novel test, the list of frameworks was evaluated for completeness with a GPT model using the prompt: “Please list all of the cybersecurity and cyber resilience frameworks” [11]. The test identified several frameworks that had not been previously identified, primarily from foreign sources such as the German government and various products from the United Kingdom.

Data Pre-Processing and Lemmatization

For both bodies of data, to enable efficient and accurate results from algorithmic analysis, the text must be preprocessed. All processing was conducted using Python 3.11.7 deployed through Anaconda Navigator with the packages as defined below. The general steps to this process include:

- Importing documents into a machine-readable format using the `pdf` package and extracting the text using the `fitz` package (also known as `PyMuPDF`) or converting comma-separated variable files using the `csv` package [12], [13], [14].
- Preprocess the text to remove stop words, URLs, and proper names or organizations. Stop words are common words in the English language that appear frequently, but do not add any value in the context of the document—the, of, or, my, etc. This step supports the subsequent lemmatization of the text by ensuring that the relevant words are retained and superfluous words are excluded [15]. Stop word removal is performed using the Natural Language Toolkit (`NLTK`) package [16]. URLs are removed with the Regular Expressions `re` package, and proper names or organizations removed with the `spaCy` package [17], [18].
- Next, the text is tokenized and lemmatized. This converts the document into a set of tokens, which can represent words or sentences. Tokenization is a technique for feature extraction within a file. Lemmatization reduces similar words into a single, meaningful, base word for analysis. For example, if there were tokens containing "running," "runs," and "ran," lemmatization reduces those three words in the tokens to "run" as the base word. The `NLTK` package carries out tokenization and lemmatization [16]. The lemmatized files are saved as both text and `pickle` files using the `os` package [19], [20].
- The `NLTK` package provides a list of common or general words, but to be effective in narrow applications, additional stop words should be added to the dictionary [16]. The lemmatized text files are assessed with the `collections` package for the top 25 words that appear in each file [21]. These words are evaluated for inclusion in the stop word

dictionary using heuristic methods to determine relevance. This analysis was repeated for several iterations to identify any additional stop words. Through this process, 66 additional stop words were added to the dictionary. The additional stop words added to the dictionary are provided in Appendix 1 for reference.

Phase One: Literature Review, Definitions, and Establishing a Coding Methodology

Objective: Through a systematic and interdisciplinary literature review, develop a coding dictionary and methodology to assess existing cybersecurity and resilience frameworks in Phases Two and Three.

Phase One begins with a broad literature review spanning multiple disciplines and uses that information to develop a coding methodology, dictionary, and research-derived scaffolding or categorization matrix for use in Phases Two and Three. Figure 4 provides a graphical view of how Phase One will proceed from literature review to a final classification scheme.

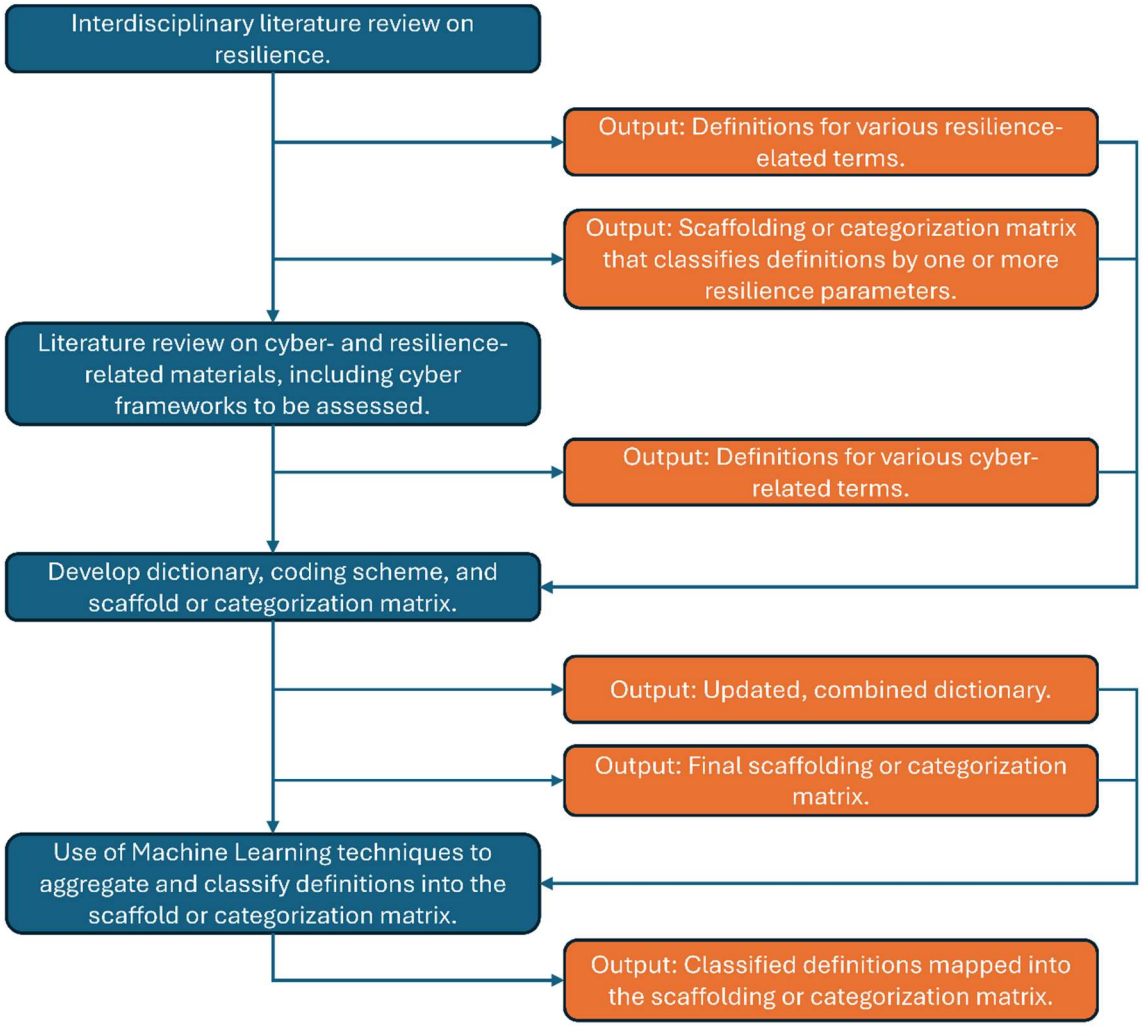

Figure 4: Flowchart depicting the progression of Phase One from literature review through a final, classified set of definitions into a research-derived scaffold.

The literature review contains three broad lines of effort. First, an interdisciplinary review of the existing body of research into complex adaptive systems, resilience, ecosystems, and their behaviors across time and scales. This review is intended to provide an integrative grounding in how the concept of resilience has evolved from its origins in ecology with other, related fields of study, such as research into high reliability organizations, safety culture, resilience engineering, organizational studies, etc., and introduce the core concepts of resilience that span across disciplines. This first line of effort will produce the necessary definitions and scaffolding for the

quantitative evaluation and assessment of the existing cybersecurity and resilience frameworks to test the three hypotheses.

Second, this study will examine the holistic attributes of cybersecurity and resilience policies from the organizational to national government and above levels. This section introduces the frameworks that will be coded and evaluated and provides a modest description of the documents juxtaposed against the derived resilience attributes. To address the aspect of scalability or extensibility in resilience, in particular, this section will provide a high-level assessment of the intended audience or target population of the policies. The frameworks will not be included in the dictionary since the cyber-related terms will come from other sources, and the frameworks would introduce bias into the analysis during model construction.

Third, the results from the prior two lines of effort will be combined with existing literature on cybersecurity and cyber resilience definitions to produce a coding scheme and dictionary. The dictionary will be coded into the scaffolding framework from the first line of effort, such that the various definitions can be assigned to an element in the matrix. For example, the definition of a firewall may be coded as both "technical," since it is a specific technical control, "organizational," since firewalls are primarily applied at the organizational level, and "machine speed" or the equivalent, to denote that this concept acts at machine, vice human, speeds, to accomplish its functions. Based on the prior research discussed in Chapter 4, it is anticipated that the scaffold will be three dimensional, with the spectrum of sociotechnical systems on one axis (e.g., technical system to sociotechnical system to ecosystem), the temporal flow of resilience (e.g., plan, absorb, recover, and adapt) on another, and the critical attributes of resilient systems (e.g., adaptive management, cross-scale interactions, thresholds, etc.) on the third.

Machine learning techniques will be applied to assist in definition aggregation and coding and provide a quantitative assessment of key definitions. Definitions are taken primarily from official sources, such as the National Institute of Standards and Technology and the Canadian government, but also, especially in the case of cyber resilience, from peer reviewed literature and from various websites from technology and cybersecurity companies discussing cyber resilience. Using similar techniques for data curation as described in an earlier section, the set of definitions will be analyzed using a BERT or Generative Pre-Trained Transformer (GPT)-based model to provide an aggregate definition of each term that is the statistical combination of the most relevant and prevalent words. These definitions will then be lemmatized for use as the baseline during framework coding and classification. The lemmatized definitions will then be classified into the scaffold using the coding scheme. The final dictionary of stop words, aggregate definitions, and categorization matrix (scaffold) is provided in Appendix 1.

A Note on Human Judgement in Phase One

Human judgment plays a critical role in Phase One, particularly in the interdisciplinary literature review and the development of the coding scheme and dictionary. While the study employs algorithmic tools to process and analyze vast amounts of text, the initial identification of key resilience concepts and the construction of a coding matrix requires human expertise. Researchers must interpret the diverse definitions of resilience from multiple disciplines—ranging from ecology to cybersecurity—and synthesize these concepts into a coherent framework. This involves subjective decisions about which resilience attributes are most relevant and how they can be mapped to the cybersecurity context. For example, resilience concepts like "adaptive capacity" or "cross-scale interactions," which originate in fields like environmental

science, must be interpreted and redefined in a way that makes sense for organizational or technical resilience in cybersecurity. This manual interpretation ensures that the coding matrix reflects a holistic understanding of resilience, grounded in expert judgment rather than solely relying on algorithmic output.

Additionally, human judgment is essential when refining the coding dictionary and adjusting the classification scheme. As definitions are aggregated and categorized using machine learning techniques, such as BERT, researchers must manually assess the accuracy and relevance of the outputs. Since algorithms can misinterpret context or conflate similar but distinct terms, human oversight is necessary to ensure that resilience attributes are correctly classified and that important nuances are not lost in the analysis. For instance, technical terms like "firewall" may be categorized under "technical controls," but human judgment is required to ensure it is also recognized as part of broader organizational resilience strategies. This iterative process of validating algorithmic outputs through expert evaluation is crucial to the accuracy and integrity of the coding matrix, ensuring that the final framework is both comprehensive and applicable to the complexities of cyber resilience.

Phase Two: Application of the Coding Methodology to Existing Cybersecurity Frameworks

Objective: Apply the coding methodology on the existing cybersecurity frameworks to allow for a structured evaluation of the frameworks against the scaffold or categorized matrix. This will produce a mapping of the contents of the frameworks in the scaffold.

Phase Two uses machine learning techniques to train a simple classification model to classify the existing cybersecurity frameworks. Figure 5 provides a graphic flowchart showing the activities of Phase Two.

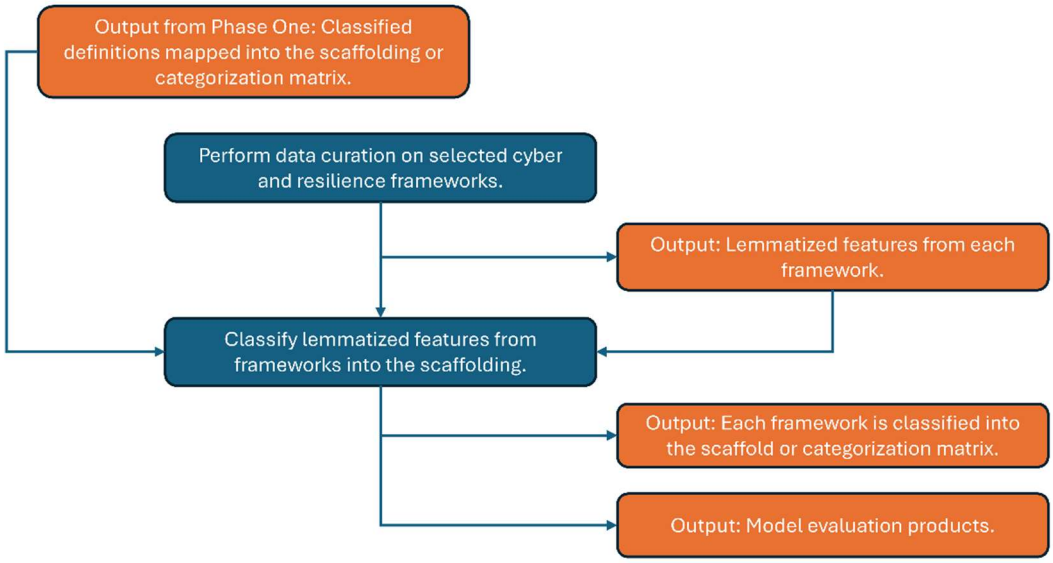

Figure 5: Flowchart of Phase Two activities, going from data curation activities to classification of frameworks into the scaffolding from Phase One, and model evaluation

The objective for this classification scheme is twofold. First, to classify the features of each framework into the scaffolding, showing which elements the framework covers, and how strong that coverage is. Second, to gather and store frequency information for term usage for use in Phase Three as part of the quantitative analysis.

This phase uses the (tf*idf) algorithm with the coding dictionary to classify features in each of the frameworks. A representative subset of the frameworks used in the analysis will be selected to train the classification model, which will then be tested against the remaining frameworks to assist in feature extraction and classification. Results from the tf*idf algorithm

will be manually assessed at several points during the process and formally evaluated using accepted techniques, described in the next section on model validation and reliability.

As a form of pseudo-control for the framework, non-cyber related resilience frameworks will be evaluated and mapped into the scaffolding to help understand the overlaps, gaps, or relationships between how the cybersecurity community interprets resilience with how other communities, such as for climate change or disaster preparedness, interpret resilience. It is anticipated that these frameworks will not map into the technically oriented categories in the scaffold (as one would expect), but will map into other domains, such as those relating to cross-scale interactions, adaptation, or state changes and thresholds.

Model Validation and Reliability

Ensuring the validity and reliability of the models used in this study is critical to producing accurate and actionable results. The machine learning models, including term frequency-inverse document frequency (tf*idf) and BERT, play a central role in classifying and analyzing the resilience attributes in cybersecurity frameworks. To validate these models, a combination of iterative, cross-validation and manual checks will be employed. Thus, the study can assess the robustness of the model and its ability to consistently classify resilience attributes across different documents.

To enhance the reliability of the model, the coding scheme and classification results will be manually reviewed at various stages. This iterative process, where the model's classifications are compared against human judgment, serves as a reliability check to identify and correct any errors or biases introduced by the algorithm. For example, during the preprocessing stage,

lemmatization and stop word removal could inadvertently eliminate important terms. By manually reviewing the most frequent terms and their classifications, the research can ensure that essential resilience-related terms are correctly captured and categorized. Any discrepancies identified in this review process will inform adjustments to the model and the coding scheme, ensuring that the classification results remain accurate and representative of the underlying data.

Finally, to ensure consistency and repeatability, the research methodology employs a standardized pipeline for data preprocessing and model evaluation. By documenting each step—from document ingestion to text preprocessing, feature extraction, and classification—the study creates a replicable framework that other researchers can follow to validate or extend the results. Reliability is further supported by using widely recognized machine learning libraries such as the Natural Language Toolkit (NLTK), BERT, and tf*idf, which have been validated in prior academic studies [22-25]. This standardization minimizes the risk of model variability and ensures that the results are not dependent on ad hoc processes or subjective decisions. Through these validation and reliability measures, the study aims to produce a robust and trustworthy analysis of cybersecurity frameworks, contributing meaningful insights into the development of a comprehensive cyber resilience framework.

Phase Three: Systematic Quantitative Evaluation of Existing Cybersecurity Frameworks and Resilience

Objective: Through statistical and other quantitative techniques, provide a systematic evaluation of the existing cybersecurity frameworks to determine if any are applicable with respect to resilience attributes (H3).

Phase three uses the results from the classification model and algorithm telemetry to conduct a quantitative analysis of the frameworks in two parts. First, each framework will be compared to the various categories that it has been assigned to, noting the strength of the classification to that category and how many features were classified into that category. It is anticipated that a technically oriented framework will not have features classified into one or more elements of the conceptual scaffolding (H3), and the language used in the framework, from a relative perspective, will be significantly more technical in nature (i.e., more use of terms like firewall, data loss prevention, malware, etc.) than resilience in nature (i.e. adaptive capacity, ecosystem, etc.) (H4). It is further anticipated that no frameworks exist that will meet most or all the interdisciplinary attributes of resilience (H2).

Phase Four: Gaps Analysis of Cyber and Resilience Frameworks

Objective: Provide a mapping of existing frameworks into the scaffolding to highlight where the existing literature may not address some or all attributes of resilience (H1).

Many cybersecurity and resilience frameworks emphasize qualitative guidelines or best practices, but few offer quantitative metrics that organizations can use to systematically measure their progress toward resilience. This integrative study will use the existing cybersecurity and resilience literature, frameworks, and other guidance to map existing metrics and other quantification systems into the scaffolding. Using heuristic methods for coding, each metric or quantitative system will be coded into the scaffolding. Heuristic methods will be used based on the smaller data set when compared to the definitional work needed for RQ1 and RQ2. Mapping of the existing metrics onto the scaffolding will accomplish two purposes. First, it will provide

the necessary data to either prove or nullify H5 since the mapping will reveal elements in the scaffold that lack any quantitative metrics. Second, it will provide the basis from which this research will build a broader quantification framework to measure all aspects of cyber resilience in the scaffold.

The gap analysis also leverages heuristic methods to code and categorize existing frameworks, comparing their applicability across different resilience attributes. By identifying gaps where resilience features are missing or insufficient, this phase will help clarify which resilience aspects—such as cross-scale interactions or thresholds—remain insufficiently addressed. This is crucial because organizations often struggle to operationalize resilience without clear and actionable principles to guide their implementation and track improvements. Then, future research can highlight areas where cybersecurity frameworks could benefit from incorporating interdisciplinary insights, drawing on fields like disaster management or complex systems theory, which may offer well-established measurement systems for resilience.

Limitations of the Research Plan

While this study seeks to provide a comprehensive framework for enhancing cyber resilience, several inherent limitations must be acknowledged that may influence the interpretation and generalizability of the findings. First, the study relies heavily on existing literature, frameworks, and publicly available documents, which may introduce limitations in scope. Since cyber resilience is a rapidly evolving field, new standards, policies, and frameworks may emerge after the data collection period that could offer additional insights or challenge the conclusions drawn from the current analysis. The selected documents represent the state of

knowledge at the time of study, and as such, the findings should be viewed as a snapshot of the current landscape rather than an exhaustive or final representation of cyber resilience strategies.

Additionally, the reliance on document analysis through machine learning algorithms such as tf*idf and BERT introduces methodological constraints. While these algorithms are effective for processing large textual datasets and identifying patterns, they are also susceptible to the quality and preprocessing of the data. In this study, text is preprocessed by removing stop words, lemmatizing terms, and filtering out extraneous content. However, errors in these steps, such as the misclassification of key terms or the removal of relevant contextual words, could lead to inaccurate results or biases in the analysis. Although manual, heuristic evaluations and iterative refinement of the coding scheme are incorporated to reduce these risks, algorithmic limitations remain a key challenge in ensuring the accuracy of results.

Another significant limitation involves the interdisciplinary nature of resilience. While this broad approach is advantageous in capturing various aspects of resilience across different fields (e.g., complex adaptive systems, ecology, psychology), the definitions and principles from these diverse domains may not perfectly align with the specific needs of cybersecurity frameworks. This introduces the possibility of conceptual misalignment, where resilience concepts from one discipline may not fully apply to another, or where the interpretation of resilience varies—though an unstated objective of this research is to show that the is little to no conceptual gap, just misapplication. This conceptual divergence could complicate the gap analysis and framework development, leading to a final framework that might not perfectly fit all organizational contexts or types of resilience challenges.

Furthermore, the study is limited by its focus on the English-language documents and sources. Given that cyber resilience is a global issue, the exclusion of non-English frameworks

and literature may result in an incomplete view of how resilience is understood and applied in different regions. Countries with distinct cybersecurity policies, especially those with emerging technologies or national security policies shaped by different regulatory landscapes, may offer unique perspectives that are not captured in this study. China and Russia, in particular, view cyberspace very differently from Western nations [26]. The language and cultural bias in data collection, therefore, represent a limitation that could affect the global applicability of the findings.

Finally, the study's approach to developing a cyber resilience framework is primarily top-down, using existing documents and frameworks to derive insights. This means the study may lack the flexibility to account for rapidly changing threat environments or novel types of cyberattacks that are not reflected in current frameworks. Although this research develops a robust methodology to analyze and quantify resilience using existing standards, it is inherently retrospective. As a result, the findings might struggle to keep pace with new and evolving cybersecurity challenges, particularly those related to emerging technologies like artificial intelligence, quantum computing, or next-generation networks. While future iterations of the framework could address these developments, the current study's reliance on historical data and existing frameworks constrains its ability to predict or mitigate novel cyber threats that affect an organization's capacity for resilience; although, the hope is that the principles of resilience will be sufficiently universal to have greater persistence and continued applicability despite the dynamism in the present and future cyber environment.

CHAPTER 3: RESILIENCE IN COMPLEX SOCIOTECHNICAL ECOSYSTEMS

What is Resilience? Struggling for a Definition

Resilience has a problem: researchers and practitioners cannot agree on what resilience is. Much of the research follows U.S. Supreme Court Justice Potter Stewart's line of thinking, "I know it when I see it" [27]. Complicating matters further, resilience is used across disparate fields ranging from ecology to psychology to engineering systems. As this chapter will show, each field discusses resilience in slightly different terms.

Definitions abound in the cyber field, but they do not bring much clarity for non-cyber professionals on how to think about cyber resilience and often overlap heavily. The publications cited below often have more than one definition for the term, muddying the waters further. To make it even harder for practitioners to understand, marketing and business development materials from cybersecurity companies often use many of the terms interchangeably.

- **Cybersecurity**: The process of protecting information by preventing, detecting, and responding to attacks [28].
- **Cyber Survivability**: The ability of warfighter systems to prevent, mitigate, recover from and adapt to adverse cyber-events that could impact mission-related functions by applying a risk-managed approach to achieve and maintain an operationally relevant risk posture throughout its life cycle [28].

- **Resilience**: The National Institute for Standards and Technology Computer Resource Center's Glossary lists ten different definitions for resilience[1], none of which match the most-used definition from the National Academy of Sciences for resilience. "The ability to prepare and plan for, absorb, recover from, and more successfully adapt to adverse events" [29].

- **Cyber Resiliency**: The ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources. Cyber resiliency is intended to enable mission or business objectives that depend on cyber resources to be achieved in a contested cyber environment [28].

- **Robustness**: The ability of an information assurance (IA) entity to operate correctly and reliably across a wide range of operational conditions, and to fail gracefully outside of that operational range [30].

- **Reliability**: The ability of a system or component to function under stated conditions for a specified period of time [28].

- **Securely Resilient**: The ability of a system to preserve a secure state despite disruption, to include the system transitions between normal and degraded modes. Securely resilient is a primary objective of systems security engineering [28].

- **Security**: Protection against intentional subversion or forced failure. A composite of four attributes – confidentiality, integrity, availability, and accountability – plus aspects of a fifth, usability, all of which have the related issue of their assurance [28].

[1] NIST Computer Security Resource Center Glossary provides a consolidation of definitions across NIST's publications. resilience - Glossary | CSRC (nist.gov)

- **Supply Chain Assurance**: Confidence that the supply chain will produce and deliver elements, processes, and information that function as expected [30].
- **Cyber Supply Chain Risk Management**: A systematic process for managing cyber supply chain risk exposures, threats, and vulnerabilities throughout the supply chain and developing risk response strategies to the risks presented by the supplier, the supplied products and services, or the supply chain [31].
- **Cyber Risk**: Risk of financial loss, operational disruption, or damage, from the failure of the digital technologies employed for informational and/or operational functions introduced to a manufacturing system via electronic means from the unauthorized access, use, disclosure, disruption, modification, or destruction of the manufacturing system [32].

The overlap among the definition is pronounced, and the community has yet to form a consensus on what cyber resilience is and how it is distinct from cybersecurity and other cyber practices and terms—resilience is almost defined as lack of successful attacks. A recent meta-survey of resilience and cyber resilience literature by Sidney Smith of the Army Research Laboratory highlighted this fact, citing eight different literature reviews that came to the same conclusions, that there is no clearly accepted definition, and that "the definition has expanded to the point of meaningless jargon" [2], [33].

Describing Resilience: A Data-Centric Approach

Getting to a commonly accepted definition has proven elusive for this interdisciplinary concept, despite the growing, interdisciplinary body of research that supports it. A meta-analysis

of the literature reveals that resilience has several common themes or features that emerge, even if a precise definition does not.

As described in the previous chapter, this research approaches resilience from an interdisciplinary, data-centric approach. It uses modern data science techniques and machine learning algorithms to analyze a corpus of resilience definitions from literature to arrive at a set of concepts or features, if not a potential common definition. This analysis uses a corpus of 44 interdisciplinary papers, books, and documents sourced from searches of the SCOPUS database and Google Scholar from February 2023 to October 2024 [2-7], [9], [34-70]. The disciplines ranged from cybersecurity, psychology, community studies, disaster relief, ecology, network science, organizational studies, international relations, sociology, and medicine. Some of the sources contained similar, qualitative analyses of resilience definitions, and thus had multiple definitions in the source. The 44 documents contained 123 definitions, of which 102 were unique.

Unlike traditional literature reviews, this research approaches these definitions from a data science perspective, using statistical and machine learning techniques to understand how this interdisciplinary community writes about resilience. To support the analyses, the 123 definitions were tokenized and lemmatized, which removes common stop words (i.e., the, of, an, etc.), and duplicate definitions removed. A dictionary of custom stop words was developed iteratively during the analysis to remove additional stop words from the corpus, which included the following words: resilience, defined, cyber, use, way, long, well, paper, manner, also, and even. The following content analyses were conducted on the 102 unique definitions: a word cloud, word co-occurrence matrix, term frequency-inverse document frequency (tf*idf) analysis, and Latent Dirichlet Allocation analysis. Each analysis provides a different insight or look into

how the interdisciplinary community identifies and defines resilience. Taken together, the body of analyses will provide a set of concepts that form a meta-definition for resilience.

Word Cloud Analysis

A word cloud is a simple visualization of term frequency in a body of text. It provides an easy visualization of which terms occur most frequently in the body of the text and allows for greater and more rapid intuitive understanding of the important concepts. The word cloud based on the 102 definitions is shown below in Figure 6. It was produced using the `WordCloud` package (v1.9.3) for Python [71].

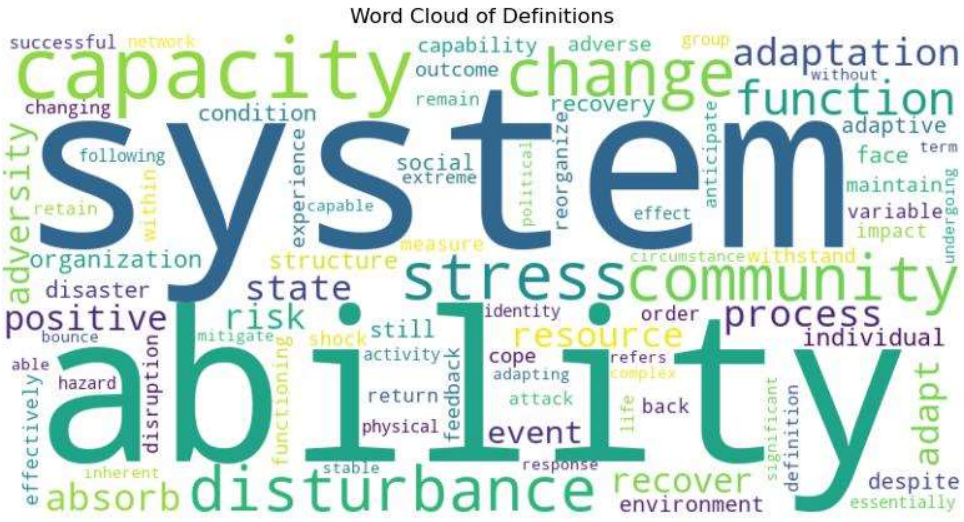

Figure 6: A word cloud showing the frequency of occurrence for terms in the corpus of 102 definitions of resilience.

From this figure, the following key concepts on resilience emerge. First, resilience is about a system's ability and capacity to change in response to a stressor or disturbance. Second, that capacity is likely a function of the system's capacity for absorption, adaptation, and community

in which the system exists. Finally, resilience is a process to recover from and overcome adversity and produce a positive state. Of note, word clouds do not show the connections, relationships, or frequency of co-occurrence between terms, just the simple counts of term frequency.

Word Co-Occurrence Network Analysis and Jaccard Similarity

Word clouds provide an excellent, first-order analysis of terms in a document, but they do not necessarily reveal linkages between the terms and how the definitions use those terms in different combinations. That linkage can be provided by a word co-occurrence matrix, which relates how often a term appears next to another term in the list of definitions. The `CountVectorizor` function in the `scikit-learn` (v1.4.2) package for Python "converts a collection of text documents into a matrix of token counts" [72]. The tokenization of the definitions decomposes each definition into a set of tokens, which are generally individual words with punctuation removed and converted to all lowercase. The function then counts how often two tokens, such as "resilience" and "system" or "system" and "ability" appear next to each other in the body of definitions.

The matrix itself can be hard to draw insights from, but viewing the results as a network diagram provides a more readily interpretable picture of how often the terms occur (node size) and how often they appear with other terms (edges connecting to other nodes). Figure 7, below, shows the network diagram for the word co-occurrence matrix on the 102 definitions.

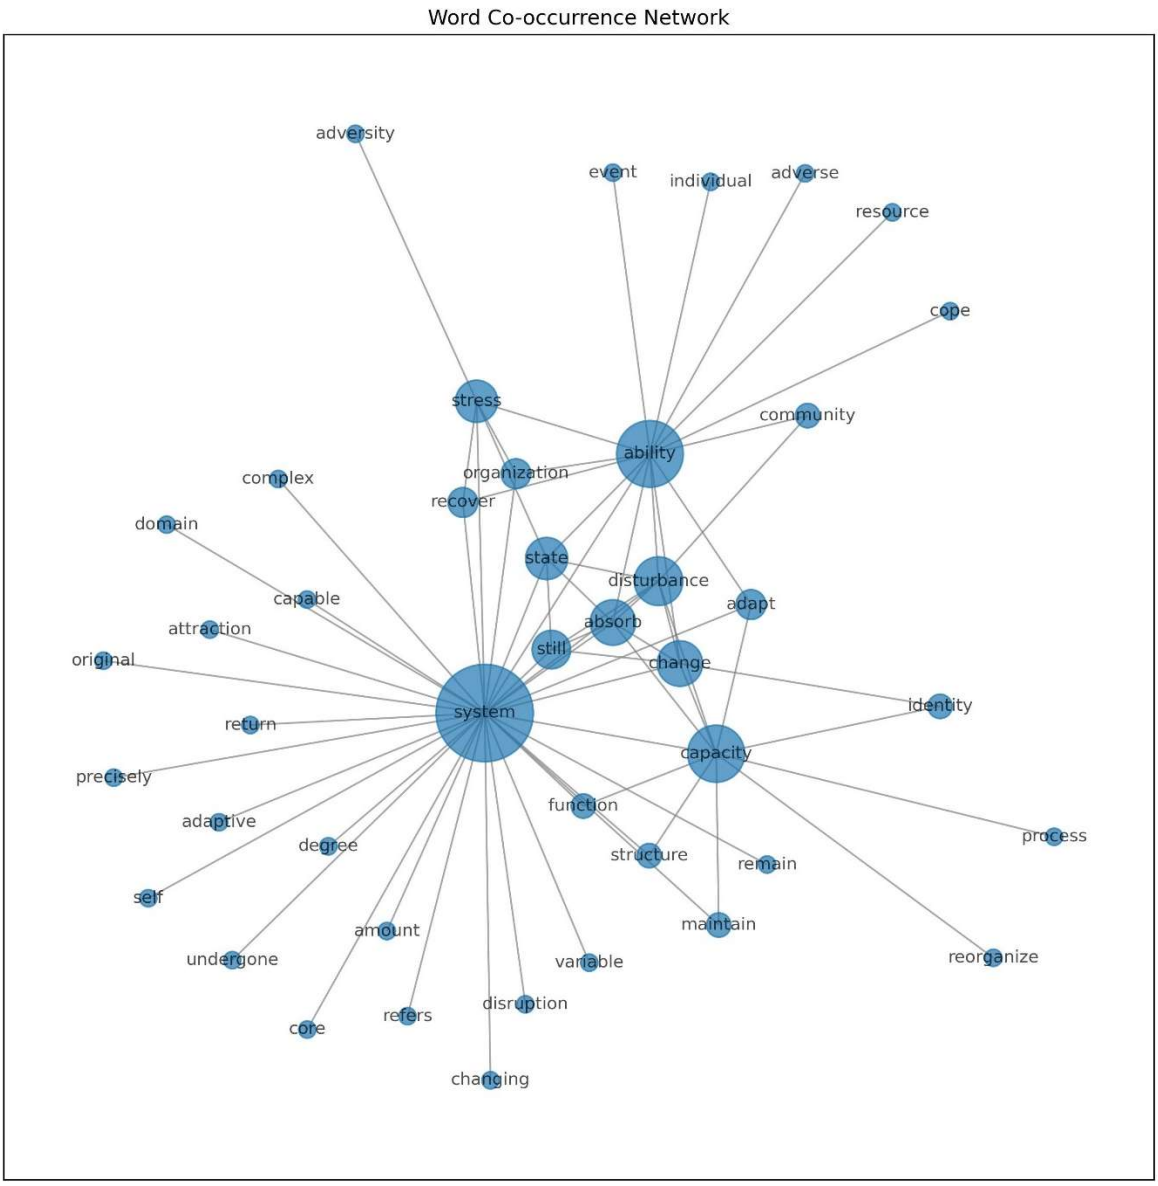

Figure 7: A Word Co-occurrence Network View of Term Frequency and Connection for 102 Definitions of Resilience

From this network diagram, resilience of a system is both a capacity and an ability. These nodes have frequent linkage to change, disturbance, state, absorption, organization, stress, and recovery. Like the interpretation of the word cloud, the word co-occurrence network diagram highlights term frequency, but also the relationship between terms.

Of greater interest might be the importance of words from the word cloud analysis and the linkage with terms in the word co-occurrence matrix. A few term groupings that are worth noting are provided here:

- System, ability, disturbance, stress, and state: Across the interdisciplinary set of definitions emerges the concept that resilience is the ability of a system to respond to a stressor or disturbance, and that system has a definable state. State can take on many connotations depending on the field, but, in general, the state can be thought of as the set of variables that define the system at a given point in time. This linkage of terms implies a temporal dimension to resilience: a system is in one state, experiences a stressor or disturbance, and responds and recovers to the previous state or adapts to a new state.
- System, structure, and capacity: The capacity of a system to respond to a stressor or disturbance is a function of its structure. This implies that there may be some system structures that reduce adaptive capacity or create outsized adaptive capacity that may not be present in similar systems undergoing similar disturbances.
- Identity, capacity, and change: Related to the previous bullet point about system, structure, and capacity, the relationship of identity to capacity and change may imply that identity plays a role in the degree of emergence, or not, of adaptive capacity in the system.
- Organization, system, stress, and ability; Community, ability, and disturbance: Both groupings relate the ability of a system to respond to a stressor or disturbance as a potential function of the community or organizational structure of system, or that the system is a part of. This implies that the system should be viewed, potentially, at multiple levels or scales.

- Capacity and process: The term capacity conjures thoughts of a static variable that describes how much of a stressor or disturbance a system can withstand while maintaining its functions or absorbing a cyberattack. The frequent pairing with process implies a far more dynamic view of capacity: that the capacity of a system may change on a temporal basis as a function of the other variables described, not as a fixed quantity from which the system can draw on or rely on.

Next, calculating the Jaccard coefficients for the pairs in the word co-occurrence network reveals insights into how similar certain words are within a set. They are equivalently known as the Jaccard similarity. The Jaccard coefficient is a function of two sets of data: the elements in the intersection of the two sets divided by the union of the two sets [73]. The coefficient is on the range [0,1], with 0 indicating no similarity between the sets, and 1 indicating the sets are identical. In this relatively small dataset, there will likely be identical sets, but a heatmap representation of the word co-occurrence by Jaccard coefficients, in Figure 8, below, shows which words share stronger relationships than the network diagram in Figure 7 leads on.

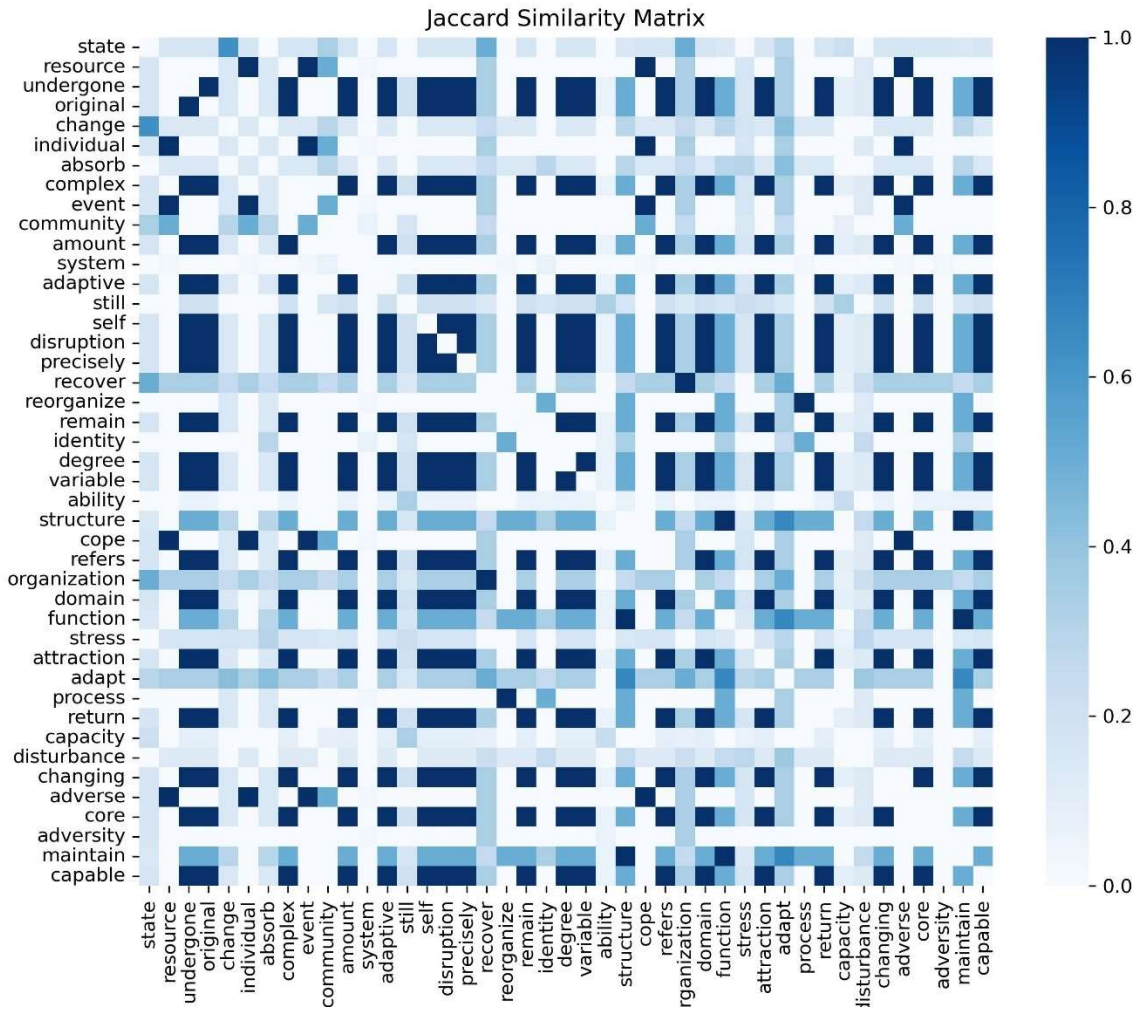

Figure 8: A heatmap of the Jaccard coefficients showing the strength of relationships between the top 29 terms in the resilience definition dataset

The heatmap confirms the previous analysis of "interesting" relationships from the network diagram. Two new insights emerge from the heatmap, though:

- Identity lacks strong similarity with most of the terms in the definition, but an analogous term, self, has strong similarity to many of the terms instead. This may be an artifact of the data analysis process in that those two definitions were not lemmatized into a single concept, but, philosophically, it raises an interesting question. Is identity tied to a sense of

self, and how does that shape the perception of resilience and what stable states are considered acceptable?

- Coping with adverse events is strongly correlated with both individuals and resources, indicating that adaptive capacity comes from both within (i.e. individual and organizational identity, as previously discussed) and the resources available to the system.

The combinatory analysis of word co-occurrence and Jaccard coefficients provided a greater level of detail into the relationships between words in the corpus of resilience definitions. It shows the nuanced and complex relationship between the system, its ability and capacity for change, and the variables that affect it. It shows that resilience is likely subjective, and has both temporal and scalable components to it, as strong ties to community showed.

Term Frequency – Inverse Document Frequency Analysis

Term frequency – inverse document frequency is a measure of a term's importance within a body of text and is a well-regarded, though basic, algorithm for analyzing text [25]. Applying the tf*idf algorithm to the corpus of 102 definitions using the `TfidfVectorizer` from the `scikit-learn` (v1.4.2) package resulted in 566 unique terms which are scored according to the tf*idf definition. The value of tf*idf analysis is the ability to reveal terms that have greater importance than others. In a collection of definitions, similar to the previous two analyses, tf*idf

highlights the terms that the interdisciplinary community uses the most. A histogram of the tf*idf scores of the 566 terms, shown in Figure 9, below, depicts this.

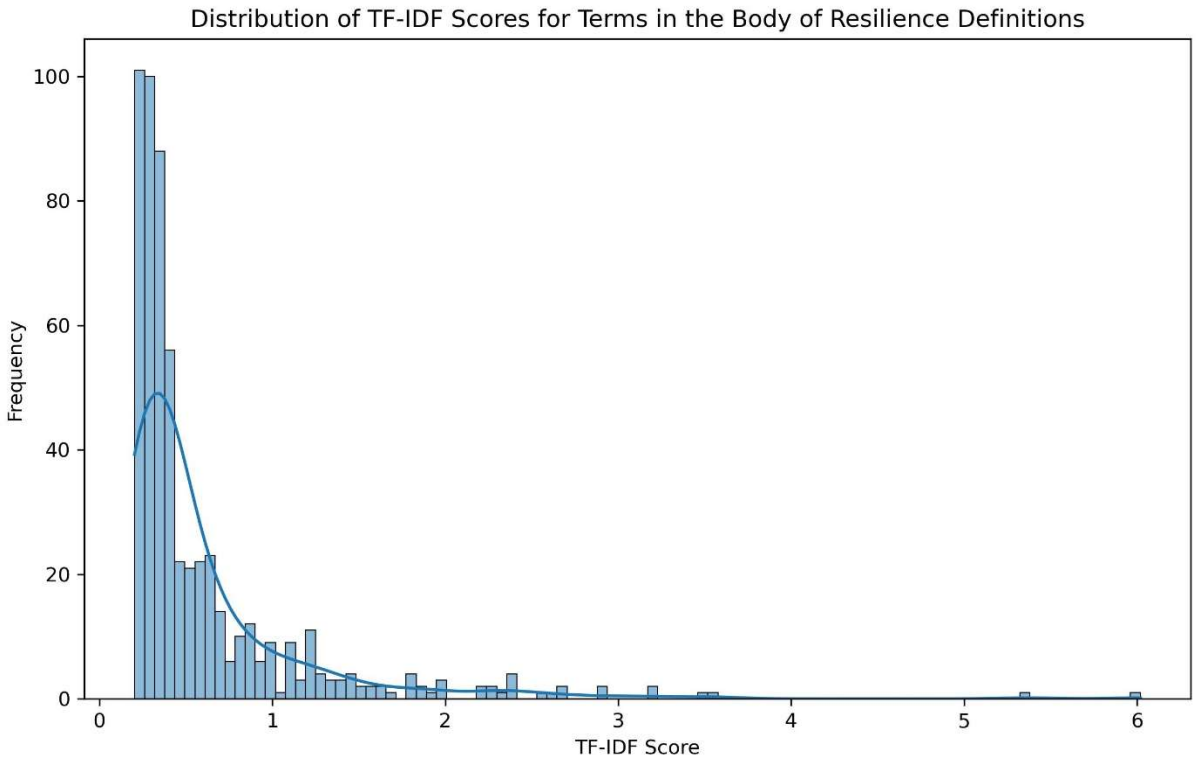

*Figure 9: A histogram of 100 bins showing the distribution of tf*idf scores for 566 terms in the corpus of resilience definitions*

The statistical description of the distribution of tf*idf scores:

*Table 1: Statistical descriptors for the histogram of tf*idf scores*

Descriptor	Value
Number of Terms	566
Mean	0.60279
Median	0.36418
Standard Deviation	0.60973

Minimum	0.20292
Maximum	6.10947

The standard practice in data science would be to examine the top quartile for the significant terms.[2] Given the number of terms and the distribution of the tf*idf scores shown in Figure 9, the top 5% of terms includes those with tf*idf scores >1.80 and provides a reasonable number of terms (n = 29) for further analysis. The top 5% of terms and their tf*idf scores are provided in Table 2, below.

*Table 2: Top 5% of terms by tf*idf Score*

Term	Td-idf Score
Ability	6.109
Capacity	5.365
Change	3.546
Disturbance	3.512
Adapt	3.212
Stress	3.195
Community	2.934
Process	2.918
Absorb	2.692
State	2.651
Recover	2.607

[2] The top quartile is the generally accepted threshold for significance in tf*idf scores; however, there is no literature directly validating this threshold, nor examining the effect of statistical distributions of tf*idf scores in selecting the significance threshold for further analysis.

Function	2.531
Risk	2.405
Positive	2.396
Adaptation	2.366
Event	2.355
Adaptive	2.312
Condition	2.267
Adversity	2.264
Resource	2.231
Face	2.212
Organization	1.970
Withstand	1.958
Adverse	1.979
Outcome	1.908
Maintain	1.879
Despite	1.862
Cope	1.827
Structure	1.811

The results from the tf*idf algorithm on the 566 terms from the definition reinforce the observations from both the word cloud analysis and the word co-occurrence matrix analysis. Resilience is most often described as an ability and capacity to change in response to a stressor

or disturbance. It is a function of the community and organizational structure in and around the system, and that adaptive capacity is a process, not a static quantity.

Latent Dirichlet Allocation Analysis

Tf*idf analysis provides an excellent lower-level analysis of a text to extract key terms, but it struggles to reduce the dimensionality of large bodies of text, making feature extraction or term significance harder to identify. Latent Dirichlet Allocation provides a generative probabilistic model to generate Bayesian topic probabilities of topics that are represented in a document [50]. By using Bayesian probabilities to estimate the probability of a term belonging to a particular topic, LDA allows for more efficient and effective topic identification and extraction.

LDA analysis requires a fixed number of topics to show the semantic similarity (coherence) among the words within the topic. Determining the right number of topics requires both quantitative and qualitative analysis. For the quantitative analysis, a grid search is performed on the dataset with a specified number of topics ranging from 5-20, which is a generally accepted adequate starting range, and the coherence scores calculated. A higher coherence score indicates greater semantic similarity within a given topic and greater interpretability of the topics. Given the very small data set this research is working with compared to other data mining and big data applications, this analysis hypothesizes that few topics will be required for a sufficient LDA analysis. Qualitatively, the topics are evaluated for consistency and overlap using heuristic methods. The qualitative analysis evaluates the topics for both meaning and distinctness from other topics. For example, too many topics would likely result in a greater degree of overlap between topics; both describing essentially or largely the

same concepts. The previous analyses—word cloud, word co-occurrence, and tf*idf—provide good calibration for the heuristic methods.

The use of coherence, log perplexity, and Jensen-Shannon Divergence provides quantitative measures to evaluate the quality of the LDA model. Coherence measures how semantically similar the words are within a topic, which corresponds well with human judgment of topic quality. Higher coherence scores indicate the topic contains words that belong together conceptually [74]. Log perplexity, the logarithm of the model's perplexity score, measures how well the model fits the data. Lower (more negative) values suggest the model better represents the underlying data distribution. However, research has shown that perplexity doesn't always align with human interpretations of topic quality [75]. Jensen-Shannon Divergence measures the distinctness between topic distributions, with higher values indicating better separation between topics. This helps ensure the LDA model isn't creating redundant topics [76]. These three metrics together provide a balanced approach to evaluating topic model quality—balancing interpretability, statistical fit, and topic distinctiveness. The results for the grid search are shown below, in Table 3, and represent the coherence score, logarithm of perplexity, and average Jensen-Shannon Divergence for the given number of topics.

Table 3: LDA grid search results on the body of resilience definitions, including model coherence, perplexity, and the average Jensen-Shannon Divergence

Number of Topics	Coherence	Log Perplexity[3]	Average Jensen-Shannon Divergence
5	0.340185	-6.50129	0.17838
6	0.313048	-6.54544	0.20177
7	0.332992	-6.56282	0.22134
8	0.360202	-6.63912	0.22890
9	0.387011	-6.60328	0.25191
10	0.379103	-6.70799	0.25381
11	0.449869	-6.78788	0.25389
12	0.396243	-6.79102	0.26114
13	0.393525	-6.76626	0.27918
14	0.377734	-6.87054	0.28048
15	0.378645	-6.93568	0.27510
16	0.347337	-6.87989	0.29107
17	0.380038	-6.93054	0.29071
18	0.386032	-6.94343	0.30051
19	0.41191	-6.90956	0.30913
20	0.431801	-7.03320	0.30037
Mean	0.37911	-6.77549	0.26111
Standard Deviation	0.03432	0.15985	0.03644

[3] The negative perplexity values are the result of the `Gensim` package (v4.3.3) taking the logarithm of the perplexity value, which is a probabilistic function. Since the probability values are less than 1, the log perplexity value will be negative.

The relatively equal scores are likely the result of a small dataset with a greater degree of similarity between terms. The maximum coherence score from the grid search, 0.449869 with z = 2.061, suggests the ideal number of topics should be set at 11. While lower perplexity values and higher Jensen-Shannon Divergence scores indicate a better model fit and greater distinctness between topics, these need to balance against the coherence score. The values for 11 topics strike an acceptable balance. 20 topics may fit better by quantitative metrics, but the qualitative analysis of the topics generated by LDA shows over-segmentation of the small dataset. The topics from the LDA algorithm are presented in Table 19, below, along with the interpretation of the grouping of terms in that topic.

Table 4: Topics and interpretations from a Latent Dirichlet Allocation analysis of resilience definitions

Topic	Terms in the Topic	Potential Interpretation
1	0.020*"ability" + 0.020*"experience" + 0.020*"stress" + 0.020*"positive" + 0.014*"system"	Resilience is related to the system's capability to handle stress and experiences in a positive light
2	0.037*"system" + 0.029*"ability" + 0.019*"process" + 0.019*"stress" + 0.013*"term"	Resilience is related to the process a system uses to respond to stressful events over a time scale
3	0.031*"system" + 0.027*"ability" +	Resilience is related to disruptive, impactful, and other extreme events and the ability to survive those events

	0.026*"disaster" + 0.026*"extreme" + 0.020*"survive"	
4	0.043*"ability" + 0.025*"capacity" + 0.017*"change" + 0.016*"system" + 0.015*"event"	Resilience is related to the ability and capacity for a system to change in response to an event
5	0.036*"ability" + 0.033*"community" + 0.021*"stress" + 0.016*"disturbance" + 0.016*"crisis"	Resilience is related to a system's functions across scales or a community (vice individual or organizational) in response to crises or stressors
6	0.054*"system" + 0.023*"capacity" + 0.022*"ability" + 0.016*"disturbance" + 0.016*"stress"	Resilience is related to a system's capability or ability to respond to a disturbance or stressor
7	0.048*"system" + 0.025*"ability" + 0.021*"within" + 0.017*"state" + 0.016*"recover"	Resilience is related to the ability of a system to recover within a given state

8	0.042*"capacity" + 0.027*"system" + 0.023*"change" + 0.021*"function" + 0.020*"absorb"	Resilience is related to a system's capacity to absorb a disturbance, continue functioning, and change in response
9	0.020*"mitigate" + 0.020*"effect" + 0.019*"social" + 0.014*"capacity" + 0.014*"environment"	Resilience is related to mitigating the effects of stressors through social capacities and the environment
10	0.037*"ability" + 0.025*"system" + 0.020*"variable" + 0.016*"positive" + 0.016*"condition"	Resilience is related to a system's ability to respond positively under variable conditions
11	0.037*"system" + 0.017*"ability" + 0.014*"shock" + 0.014*"particular" + 0.014*"adapt"	Resilience is related to a system's ability to adapt in response to specified or particular shocks

The topical distribution from the LDA analysis reveals how nuanced resilience is. Heuristically, this checks against the inability to arrive at a common definition within a single

discipline, let alone in the broader interdisciplinary community. The LDA analysis confirms the results of the prior analyses that resilience is dynamic, impacted by scale and time, and is both process and capacity.

Cumulative Results of the Data Analysis

The analysis yields interesting insights into the generally accepted attributes of resilience, and how those core attributes shape our perceptions of resilience and resilient systems. There is general agreement that resilience is a defined system moving through phases before and in response to a disruptive event, and that the quality of that reaction and response is determined by the pre-existing adaptive capacity before the event, the resources and processes available in response to the event, and how the system responds and adapts across several scales as a result. This fits with resilience as an emergent property of a complex adaptive system [77].

First, there is broad agreement that resilience is not a singular, static concept with a tidy definition, but rather a dynamic concept that changes across time and scale for a given system. Prior to a disruptive event, resilience largely revolves around the adaptive capacity of the system in question, such as an individual, community, organization, or ecosystem. This capacity is determined from the adaptive memory of responses and outcomes from earlier stressors (i.e., "lived experience" from a psychological perspective), changes to the system since the earlier stressor, and the capacities of the other actors above and below the scale of the system in question.

After a disruptive event occurs, the magnitude of resilience of the actor is determined by that pre-existing capacity for adaptation, and the processes by which the actor absorbs, responds,

and adapts to the disruptive event. This can be best seen graphically by adapting a common interpretation of resilience of plan and prepare, absorb, recover, and adapt from a disruptive event and indicating how these concepts relate to resilience as adaptive capacity and process [78].

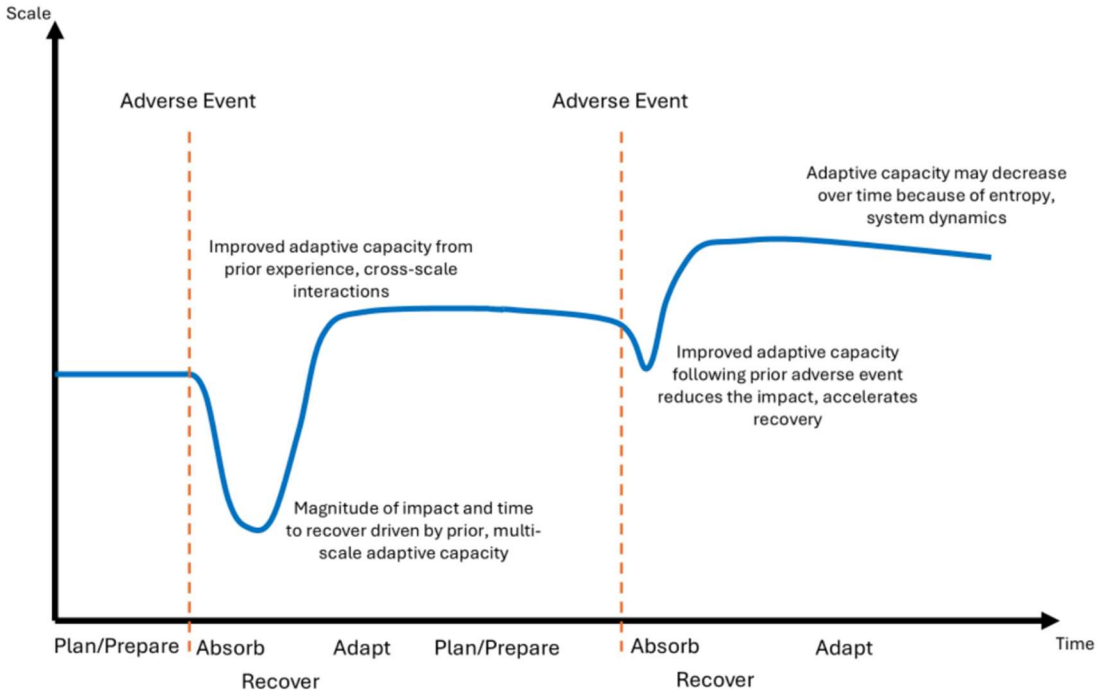

Figure 10: Resilience and adaptive capacity across time and scale in response to multiple adverse events.

Time and scale become central concepts in the pre- and post-disruptive event phases and define the flow of resilience. Establishing these as a two-dimensional plot allows for the existing interdisciplinary research to be placed into a broader or more systemic context, thus giving resilience a stronger notion of being a fundamental concept on how systems respond to a disruptive event, regardless of the topical approach (i.e., ecology, psychology, etc.). Figure 11,

below, shows this framework with the notional placement of different disciplines involved in resilience. Both axes, time and scale, are considered at log-scale to reinforce the concept of linearization about a specific point in time and scale. It also shows the intractability of thinking about resilience as a purely linear, finite concept that is disconnected from cross-scale and long-term impacts.

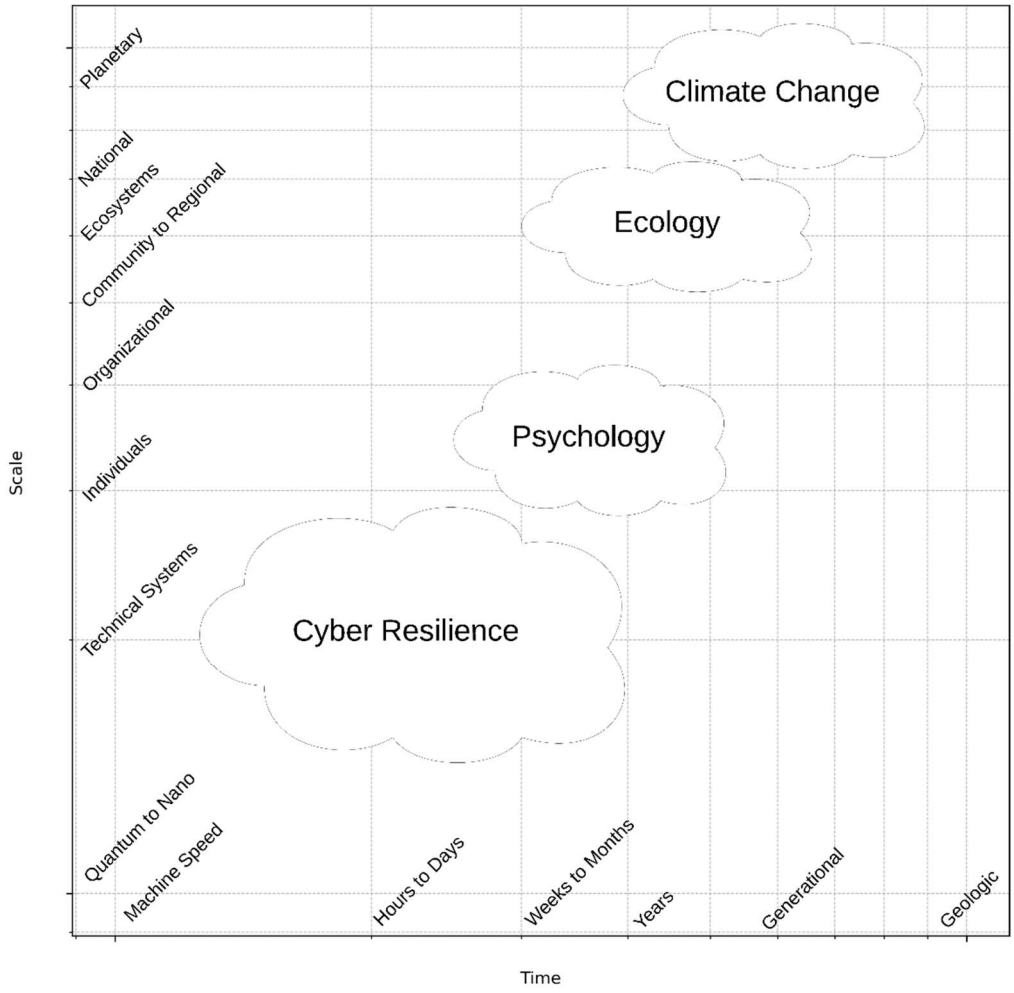

Figure 11: Resilience as a function of time and scale with notional overlay of where various disciplines linearize.

Resilience and Time

The response and adaptation phases operate at short- mid-, and long-term time scales. Short-term actions focus on immediate responses and adaptations to the disruptive event(s). For example, for a cyber breach, this includes traditional incident response actions and post-incident recovery and adaptation actions, such as implementing additional security controls, incorporating new security features, educating users, etc. The mid-term actions in response to a cyber breach might focus on improving the enterprise architecture, organizational investments necessary to reduce the risk of a future breach, or changes in policy and governance. In the case of an organization that is part of a nation's broader critical infrastructure or nested within a broader ecosystem (as all are), the long-term actions may also include systemic responses to the breach, such as changes in government policies or laws, investigatory actions, punitive actions, etc. To which the actor must adapt. For example, Storm-0558's breach of Microsoft in 2023 resulted in one of the first high profile investigations by the nascent Cyber Safety Review Board, whose procedures, actions, and recommendations will drive Microsoft and other organizations to adapt their policies to this new investigatory paradigm [79]. Though the future role of federal government oversight and investigations is now up in the air after the Trump Administration disbanded the Cyber Safety Review Board [80].

Resilience at Scale

Second, the conceptualization of a system as described in the literature is one of a set of actors, processes, and resources across multiple scales. While that can theoretically drive analysis from the nano- to galactic-scales, drawing a system boundary provides necessary constraints to the analysis. These complex adaptive systems are inherently non-linear, and defining a system

boundary is akin to assuming linearity at a specified point for ease of analysis with minimal impact to understanding (accuracy). Thus, resilience is defined at a scale and level meaningful to the point of reference, such as a psychologist working with an individual or group, an ecologist studying a local watershed, or a government agency seeking to improve resilience in a critical infrastructure sector. Brief examples from several fields highlight the benefit of a "linear" system boundary in studying resilience.

Psychological research into personal resilience focuses not solely on the individual, but how the community around that individual creates the capacity, or lack thereof, for resilience in the individual through experiences in the community and the resources of the community in helping them process a disruptive experience.

At the community level, research identifies pre-existing community attributes as central to the adaptive capacity before a disaster or other stressing event strikes, and the resources available to the community at scales above (county, state, and federal governments, non-governmental organizations, etc.), below (individuals and families, etc.), and at the same level (e.g., organizations within the community such as religious organizations, community groups, etc.) as the community [54], [56], [81].

At the organizational level, there is little direct research into resilience; however, the literature around high reliability organizations, resilience engineering, and similar concepts correlate well with the aspects of resilience discussed thus far [70], [82], [83]. There is a strong parallel to community level concepts and the availability of organizational and community resources before and after adverse events. Within the organization, the research highlights the relationships between technical systems, processes, and organizational culture—a sociotechnical

system—as key for sustaining high performance over time and adapting to adverse events in a positive fashion [55], [82-88].

At the ecosystem level, ecologists examine the resilience of ecosystems through disruptive events, and the impact of those events on ecosystem services. Research at this level is highly variable, from modeling the resilience of certain wildlife populations in predator-prey relationships to ecosystem response following fire to the degradation of ecosystem services from human impact [89-94]. Many of these studies take place at the local to regional scales, but in showing how scale and time drive ecosystem resilience, there is additional research into the impacts of climate change on local and global ecosystems as well [95].

Common Concepts from Resilience Literature

The analysis using data science techniques yielded multiple insights into the definition of resilience, but it did not fully capture several of the concepts that are still emerging, likely due to lack of sufficient presence within the definitional data set to rise quantitatively to the top. These concepts link broadly with several topics from the Latent Dirichlet Allocation analysis and are worth greater study here: adaptive capacity, graceful extensibility, adaptive management, critical functions, cross-scale interactions, and panarchy (a nested set of adaptive cycles) and the adaptive cycle. These concepts have been largely advanced by Woods, Fath, Connelly, Holling, and others, and heuristically have strong ties to the other disciplines, but the other disciplines have not yet adopted the terminology [43], [96-99].

Adaptive Capacity and Graceful Extensibility

In the realm of complex systems that serve human purposes, adaptive capacity and graceful extensibility emerge as fundamental concepts that enable sustained adaptability. Adaptive capacity refers to a system's ability to continue adapting to changing environments, stakeholders, demands, contexts, and constraints [69], [100]. Graceful extensibility, on the other hand, is the opposite of brittleness—a sudden collapse or failure when events push a system beyond its boundaries for handling changing disturbances and variations. A system with graceful extensibility can extend its capacity to adapt when surprise events challenge its boundaries [69].

These concepts are interconnected, as adaptive memory and adaptive management play significant roles in shaping a system's adaptive capacity. Previous experiences, such as coping with cyberattacks or other disruptive events, influence how individuals, organizations, and systems respond to future challenges. The response to a disruptive event is a manifestation of the adaptive capacity of the affected actors, with the processes by which they absorb, respond, and adapt determining the magnitude of resilience.

As time and scale are crucial factors in understanding both adaptive capacity and graceful extensibility, these concepts can be visualized as a two-dimensional plot that places interdisciplinary research into a broader or more systemic context (Figure 10 and Figure 11). By addressing adaptive capacity and graceful extensibility, we can better understand how complex systems sustain adaptability during changing conditions, ultimately helping us design systems capable of withstanding various challenges and adapting in the face of uncertainty [69].

Adaptive Management

Memory and adaptive management capture the fact that individuals and organizations have memory, and this influences adaptive capacity. Previous cyberattacks and other disruptive events color how individuals and the organization respond to the next one. Indeed, the response to Midnight Blizzard's cyberattack against Microsoft coming so quickly following the Storm-0558 breach has amplified the calls for greater accountability for Microsoft and driven portions of the federal government to accelerate diversification efforts [101], [102], [103]. The changes made following an attack to make the next one less impactful reflect of the adaptive capacity of the people and the organization as a whole [40]. This is firmly a social phenomenon, not a matter of technical controls and system memory, and thus brings in the social impacts from various cognitive biases in management actions with respect to preparing for the next attack. These biases shape how people respond and adapt following an attack, and the social aspects of this shape changes to the technical and sociotechnical systems [40], [104], [105]. Finally, there is a strong linkage with cross-scale interactions and adaptive management. As organizations come through particularly difficult attacks, understanding the memory and adaptive management of other organizations in the ecosystem, such as state and federal governments, is prudent.

Critical Functions

Critical functions are somewhat self-explanatory. They are the reasons that an organization exists and the value it seeks to deliver as defined by key stakeholders. These stakeholders, in identifying critical functions as they relate to organizational resilience, must identify "the resilience of what, to what, and for whom" [40]. Key to defining critical functions or services is defining both the scale of interest and the time span of interest. It is important to

note that the critical functions are not specifically cyber- or IT-related. These are *organizational* functions critical to delivering the core value proposition. IT exists to enable organizations to better deliver value and services, and cybersecurity and cyber resilience are the means to ensure that the organization can continue to deliver those services despite setbacks. These critical functions provide the basis for risk management activities, and defining acceptable levels of degradation or loss is a crucial responsibility of an organization's senior leadership [106].

Cross-Scale Interactions

Cross-scale interactions acknowledge that cyber-attacks or other disruptions, such as weather, impact the organization at multiple levels and on multiple time scales [40], [66]. The long history of cyberattacks against Microsoft provides a clear example of this [107]. Microsoft responded to each cyberattack as it came in over the years, but the ongoing perception of security flaws in their products that the cyberattacks highlight have pushed new sectors and organizations, such as the federal government, to diversify services into additional vendors–also a response to cyberattacks, albeit on a longer timescale [101], [108]. While it was likely well beyond drafting and into interagency coordination during the recent Midnight Blizzard and Storm-0558 cyberattacks on Microsoft, NSM-22 can potentially be seen as a high-level policy response that will alter how ecosystems, and the organizations in them, respond to these cyberattacks in the future. Indeed, calls are growing for the federal government to take a firmer response to Microsoft for their software security, and more broadly to the technology giants that run the bulk of the public cloud services [101], [109]. This change of both scale and temporal spans as a result of cross-scale interactions is known as *panarchy*, which captures the dynamic and hierarchical structuring of complex systems, and the cycles they go through [97].

Thresholds

Thresholds invoke the idea that a system or organization may have several stable states that it can operate in. Business continuity plans address this in more concrete terms and capture the threshold concept well. For example, a major power outage or natural disaster in one location may result in the activation of an alternate location from which a business restarts some or all critical functions. Understanding the potential stable states of an organization and how sensitive or robust they are to disturbances is critical to developing a resilience strategy [40], [67]. In the case of a federal agency facing potential service outages from a cyberattack against Microsoft Exchange services, the agency may determine back up means to ensure that a portion of the agency's services or throughput shift to alternative means while email is inaccessible, such as reducing operations to only certain critical functions, which get printed off and hand-walked through an organization if necessary, instead of all routing done by email. This could be sustained for a long period of time because the reduced number of products matches the inefficiencies from manual routing, or the organization and slowly adapt and create new processes to restore stopped workflows during the Microsoft Exchange outage. In the case of Microsoft, internally, reporting from ProPublica revealed that the Microsoft Security Response Center had crossed a threshold into a new stable state: from responding appropriately to reported threats and working with product teams to remediate them to adopting the attitude "How can I get to won't fix?" [110].

Panarchy and the Adaptive Cycle

The 2002 introduction of panarchy and adaptive cycles in ecosystems brought about a major change in the way ecologists approached the study of resilience in ecosystems [97]. The adaptive cycle linked several elements present in ecosystems–system potential, connectedness, and resilience–to how they vary over time. It defined four phases that ecosystems, or in our case, organizations, go through:

1. Exploitation and growth,
2. Equilibrium or conservation of the status quo ante,
3. Release, collapse, and the crisis,
4. Confusion, reorientation, and innovation [43], [97].

The phases of a resilience shown in Figure 10 occur in Figure 12, below, at the crisis point when an adverse event occurs that impacts the organization's ability to operate, through the confusion and innovation phases as the organization absorbs and recovers. The adaptive phase begins as the organization enters a phase of new growth, and the transition back to planning and preparation for adverse events begins anew in the new status quo.

The private sector sees this cycle play out all the time in the sensational stories of the rise and fall of organizations, turnarounds, and companies beating the odds to stay on top. Think of *Good to Great* by Jim Collins, *The High Velocity Edge* by Steve Spear, and countless memoirs from business leaders telling the story of their company's success [111], [112]. These stories all have a common thread. Organizations accrue resources, including people, financial assets, intellectual property, etc., and then leverage these resources for tremendous growth. Once in a satisfactory position, whether in terms of product, market share, or myriad other indicators, they

seek to conserve their position, which inevitably faces a crisis. That crisis results in a collapse or lesser release of resources, which prompts a reorganization, realignment, or other initiative to reposition the organization for a new phase of growth. Figure 12 gives a visual representation of how this occurs over time, as adapted for social systems.

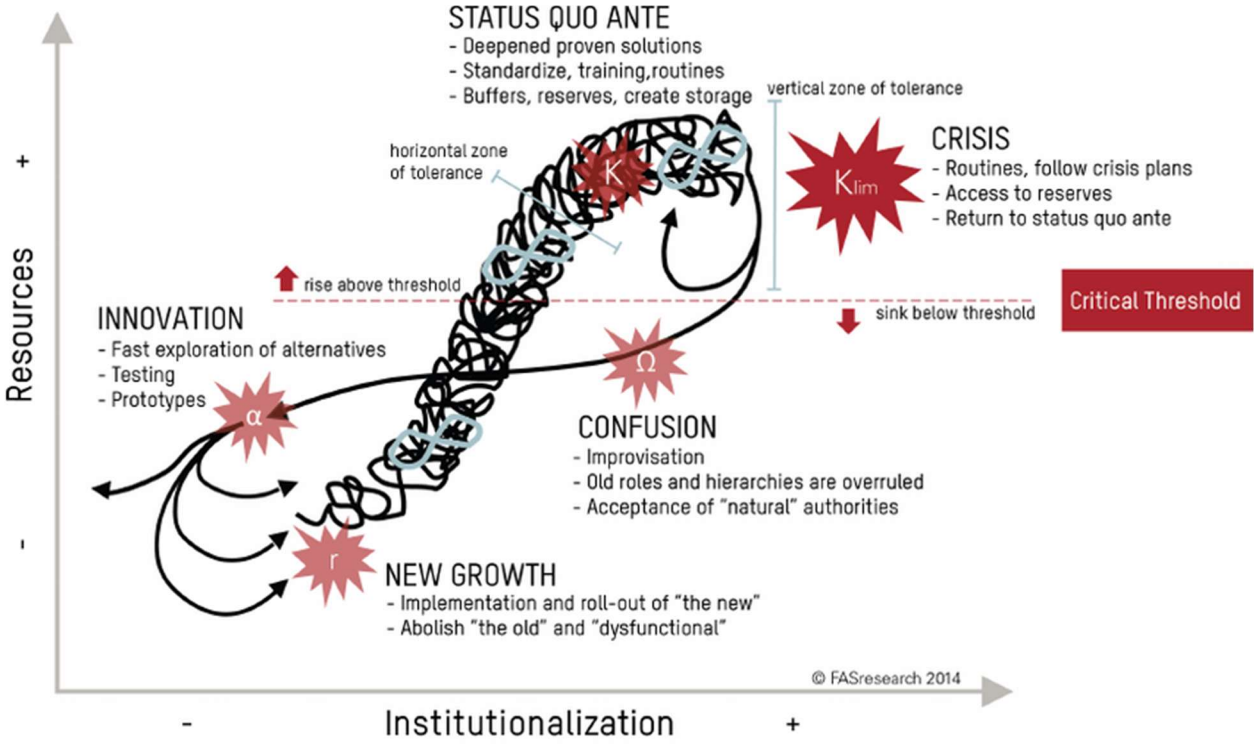

Figure 12: The adaptive cycle as applied to social systems from [43].[4]

Organizations spend most of their time in the exploitation and conservation phases. How long the organization lasts in those phases depends on how resilient and adaptive the organization is to threats, changes to the industry, and myriad other disruptions. The internet unicorns of the dot com bubble seemed to rise and fall overnight while the venerable General

[4] "Figure 3: Adaptive cycle applied to social systems. Stages in this cycle are similar to ecological stages, from new growth to status quo, to confusion, and innovation. The differentiation between crises that remain within the threshold and those that lead to dissolution are indicated by the vertical range of tolerance." Fath BD, Dean CA, Katzmair H. Navigating the adaptive cycle: an approach to managing the resilience of social systems. *Ecology and Society* 2015;20. Figure used with permission from the author under the Creative Commons License 4.0 - https://creativecommons.org/licenses/by/4.0/.

Electric managed to sustain a minimum level of performance for more than 120 years before being delisted from the Dow in 2018 [113]. In the case of Microsoft's security culture, we can "watch" the cycles play out over the last twenty years, first in 2002 with the secure development lifecycle and Trustworthy Computing initiative, to Bill Gates' memo to the company in 2012 redefining and reinforcing Trustworthy Computing, to Chief Executive Officer (CEO) Satya Nadella's response to the Storm-0558 breach with the Secure Future Computing Initiative [114], [115], [116]. Each new pronouncement punctuates the crisis point and potentially begins the period of confusion as Microsoft begins to reorganize and reprioritize their work in response.

Mapping the adaptive cycle onto an organization to understand how its performance evolves over time is not well established, nor does it provide individuals with an understanding of what drives the seemingly random movement of the exploitation phase. Borrowing a concept from safety-critical industries like nuclear power and healthcare, Cook and Rasmussen provide a framework in Figure 13 that shows how the operating envelope that encompasses the variations in organizational performance and how changes in the organization can push it into a state of crisis, confusion, or collapse as it crosses an unacceptable boundary.

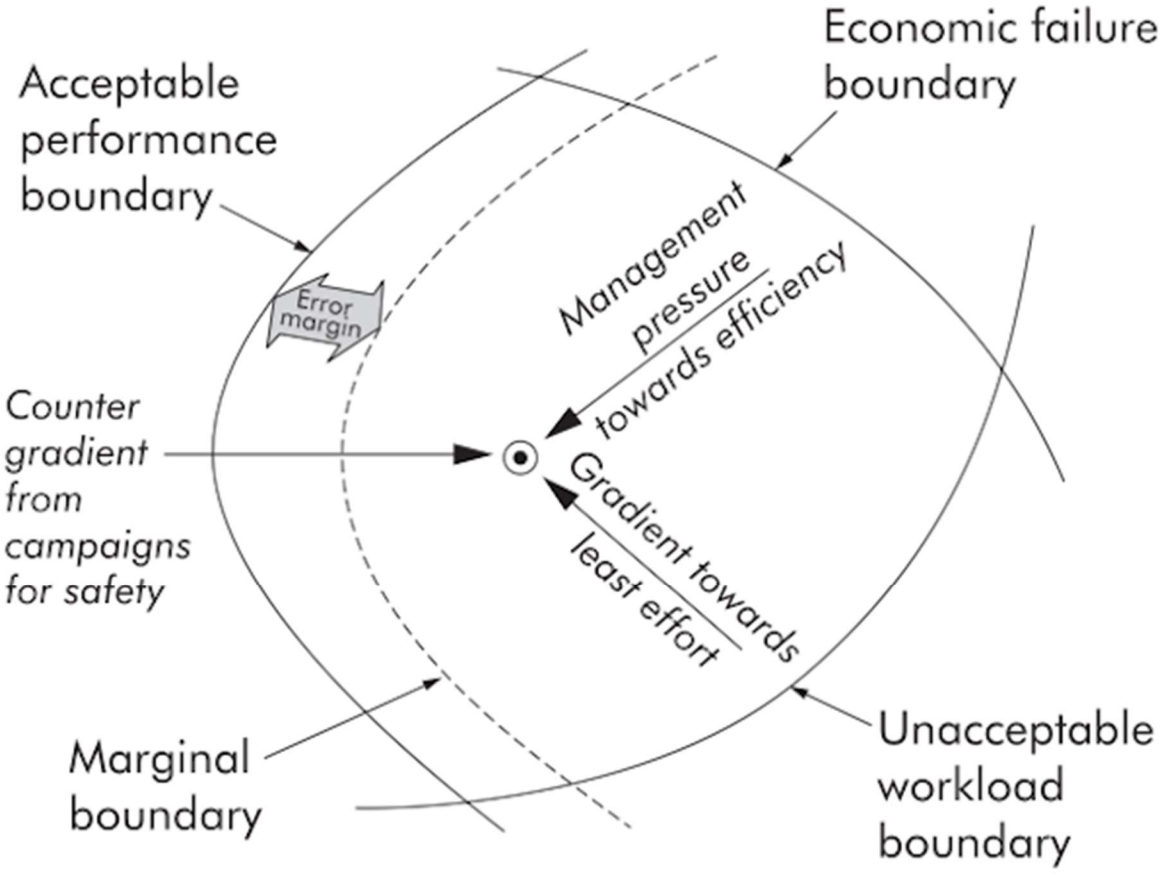

Figure 13: The dynamic safety model from [117] which shows how changes in the forces acting on an organization can drive an organization outside of acceptable performance.[5]

From this, the random wanderings shown in Figure 12 begin to make more sense from an organizational perspective. As people come and go, culture changes, the organization responds to shifts in the industry or threat environment, etc., the performance of the organization wanders in this envelope. In the years leading up to the Midnight Blizzard breach, we can theoretically "visualize" how Microsoft writ large and the Microsoft Security Response Center move

[5] "Figure 1: Dynamic safety model. (A) Gradients push the system operating point away from the boundaries of economic failure and work overload and towards the unacceptable performance (accident) boundary. (B) Stable low risk systems (A) operate far from this boundary; stable high-risk systems (B) operate nearer the acceptable margin but the operating point moves in small increments and remains largely inside the marginal boundary; unstable systems (C) have large rapid shifts in the operating point." Cook R, Rasmussen J. "Going solid": a model of system dynamics and consequences for patient safety. *BMJ Quality & Safety* 2005;**14**:130–4, by permission of BMJ Journals. This image is not covered by the terms of the Creative Commons license of this publication. For permission to reuse, please contact the rights holder.

behaviorally from an unacceptable workload boundary because of the number of bugs being reported along with management prioritizing the shipping of new features and the Azure cloud products until Microsoft crosses a final, acceptable performance boundary that enables Midnight Blizzard to make it into the networks [79], [110].

General versus Specified or Narrow Resilience

The relationship between resilience and the disruptive event requires elaboration. Both academic literature and popular discourse on resilience often fail to specify the nature of the relationship, which hides a critical aspect of resilience. The vast majority of definitions do not explicitly define this, but a sufficient number did reference it such that the 11th topic in the LDA analysis produced the key term "particular." General resilience shares similar aspects with the debate around artificial general intelligence and, more simply, optimization problems.[6] Achieving optimality for all cases generally results in suboptimal results for all case. No system can be optimal, or resilient, to all conditions or threats at all times. The law of entropy ensures that even if perfect optimality or resilience is achieved, it cannot be maintained. The cross-scale dynamics, emergent properties resulting from human involvement, and other factors ensure that entropy acts on the system.

[6] Of note, optimization is sued here in the metaphorical sense, not the literal sense. Adaptive capacity, and thus resilience, by definition, will be suboptimal since greater resources, capacity, people, etc. will be required to achieve resilience than the results of an optimization study would prescribe. Optimization drives for the most efficient use of resources, and that efficiency can actually result in the system becoming less resilient over time as the adaptive capacity is further removed from the system, beginning at the technical levels and move through the sociotechnical and higher levels as the dynamics of the system change in response to the optimization goals.

CHAPTER 4: SYSTEMATIC EVALUATION OF EXISTING CYBER GUIDANCE AND FRAMEWORKS

Similar to the development of resilience definitions in the previous chapter, evaluating cyber guidance requires a multi-source and multi-faceted approach. Understanding the co-occurrence of resilience terms with cyber terms is critical to determining whether the existing cyber guidance and frameworks have sufficient scope on the time and scale aspects of resilience. This chapter includes three major processes: 1) development of a cyber and resilience dictionary, 2) classification of dictionary terms by time and scale, and 3) evaluation of existing cyber guidance and frameworks against the dictionary and previous results from the resilience analysis in the last chapter.

The development of the broader cyber dictionary proceeds as follows. First, the definitions will be cleaned, tokenized, and lemmatized for follow on analysis. Second, based on the volume of definitions present in some of the sources, the corpus will be evaluated for down selection using BERT algorithms to determine the most relevant or popular terms. Third, the curated dictionary will be evaluated for existing resilience terms, and, if lacking, the terminology will be added to provide for follow on co-occurrence analysis.

Second, the terms in the cyber dictionary are classified by time and scale based on the concepts developed in the previous chapter to develop a two-dimensional scaffold on which the existing cyber guidance, standards, and frameworks can be evaluated. The objective of this scaffold is to show how the distribution of cyber and resilience terms is distributed across time and scale alongside the co-occurrence between cyber and resilience terms.

Lastly, a broad search for existing cyber and resilience standards, guidance, and frameworks is curated for analysis. The data set primarily includes guidance for cybersecurity processes and practices but will also be expanded to include resilience guidance for other areas, such as climate change, critical infrastructure, and national resilience to help baseline the results from the analysis in this chapter. These existing cyber guidance, standards, and frameworks will be assessed using both the scaffold and the prior data-centric analyses of resilience definitions. This process will elicit the time and scale aspects of resilience in the existing guidance and provides a parallel analysis of co-occurrence and cosine similarities for resilience-oriented terms and topics from the tf*idf and LDA analyses in the previous chapter.

Curating a Dictionary of Cyber Definition using Large Language Models

Several government and industry sources provide extensive and curated lists of cyber definitions. 4,465 terms were scraped or downloaded from the National Institute of Standards and Technology Computer Security Resource Center, the Canadian Center for Cyber Security, the National Initiative for Cybersecurity Careers and Studies program at the Cybersecurity & Infrastructure Security Agency, the Committee on National Security Systems, and the SANS Institute [29], [30], [118], [119], [120]. The definition set contains 9,203 definitions for the 4,465 terms, indicating duplicate coverage for most terms, at a minimum. The statistical descriptors for the number of definitions per term are shown in Table 5:

Table 5: Statistical Descriptors for Curated Dictionary of Cyber Terms

Descriptor	Value
Total Number of Terms	4,465
Total Number of Definitions	9,203
Minimum Number of Definitions of a Term	1
Maximum Number of Definitions of a Term	121
Average Number of Definitions of a Term	2.061
Standard Deviation of Number of Definitions per Term	3.285

The term "resilience" set the maximum number of definitions at 121, as described in the previous chapter, since those definitions were added to the list of terms. Without that term, the top ten terms with the most definitions are: assessment, cryptographic key, information system, user, authenticate, cyber threat, key, access control, digital signature, and approved (relating to approved cryptographic algorithms), which had between 24-43 definitions each.

Assessing Definitions using BERT Algorithms

Working with 4,465 terms in assessing the existing cyber guidance would be cumbersome, and many of the definitions have narrow applications in cyber fields, such as three-key triple data encryption algorithm (3TDEA), derived relationship mapping, prime number, or s-box. To make data management and analysis easier with minimal loss of fidelity, the next phase evaluates the dictionary terms using BERT algorithms for popularity or relevance. The use of BERT algorithms allows the corpora of training data and weights in those algorithms to be used to determine relevance or popularity of the term within the broader corpus, whether a general

training data set or a specified and tuned data set. For this analysis, a dual evaluation will be performed using the DistilBERT Base Uncased algorithm, fine-tuned on the Stanford Sentiment Treebank (SST-2) dataset, and the SecureBERT Plus algorithm, a RoBERTa algorithm fine-tuned on 253,433 cybersecurity documents [121], [122]. Both algorithms were sourced from Hugging Face and implemented using the following python packages: `json` (v3.13), `matplotlib` (v3.8), `Seaborn` (v0.12.2), `transformers` (v4.42.3), `pandas` (v2.1.4), `NumPy` (v1.24.3), and `Scikit-Learn` (v1.4.2) [123-131]. DistilBERT provides a faster and more computationally efficient BERT algorithm that has been trained on a broad or general dataset. It provides an appropriate counterweight to a tuned cybersecurity algorithm like SecureBERT Plus to prevent biasing the results excessively toward the existing cyber literature. A weighted average output of the two algorithms will create the final dictionary of terms.

The two models provided different results, as expected, based on their training corpora and colloquial usage in those corpora. The DistilBERT algorithm produced the following results, visualized in a word cloud and as a histogram by term relevance, for the most relevant terms from the dictionary in common usage, shown in Figure 14 and Figure 15, below.

Figure 14: Word cloud of top 500 terms from DistilBERT ranking of cyber terms.

Figure 15: Histogram of cyber term relevance using the DistilBERT algorithm.

A qualitative review of the top 500 terms explains the left skew of the histogram: most of the terms are not niche cyber terms, but terms that have multiple meanings within the broader English language.

The SecureBERT Plus algorithm, on the other hand, produces a histogram that is closer to normally distributed rather than the extreme skew of the DistilBERT results. This is an artifact of the cyber-centric training corpora of the model. Since taking the top 500 results by relevance score will only select the right-hand tail of the distribution, the terms within ±0.25 standard deviations of the mean score will be selected. This selection was manually verified by checking for common words, such as information, firewall, access control, etc. Those common terms were not in the selection from the right hand tail but were present in the mean selection. This approach yielded 939 terms, which are highlighted in the histogram showing in Figure 16 below.

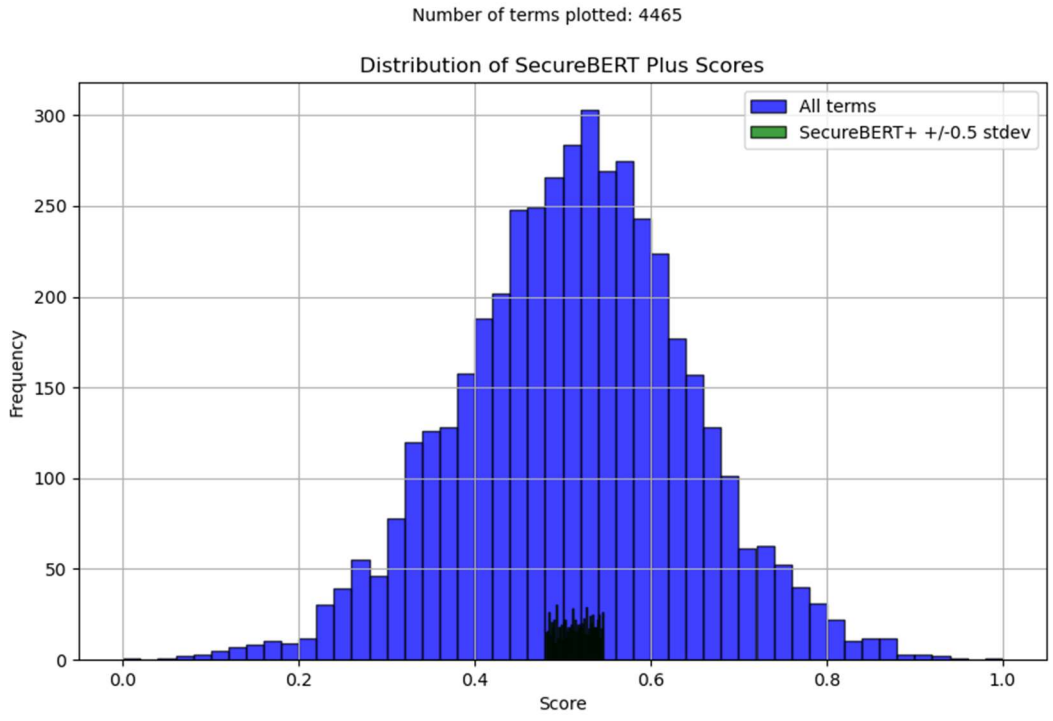

Figure 16: Histogram of cyber term relevance using the SecureBERT Plus algorithm

The terms in this set are visualized with a word cloud, shown below in Figure 17.

Figure 17: Word cloud of terms within +/- 0.25 standard deviations of the mean score from the SecureBERT Plus algorithm

The terms selected from the outputs of these two models do have some overlap, as shown in the Venn diagram in Figure 18. The diagram is based on the outputs of both models, not the weighted average that will be performed to develop the final dictionary.

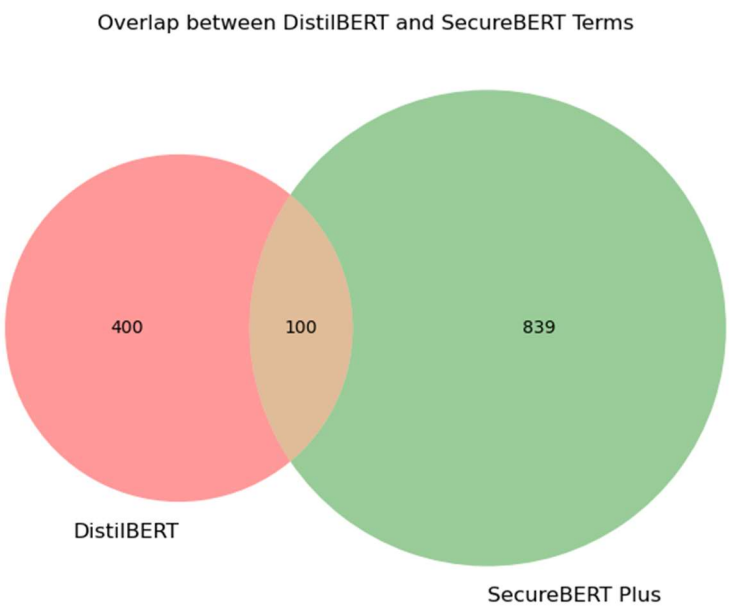

Figure 18: Venn diagram showing overlap of the two sets of top terms from the DistilBERT and SecureBERT Plus models.

The combined term set that will form the dictionary is combined by weighing the scores for a term from the DistilBERT model and the SecureBERT Plus model in a 25/75 split. Selection of the weighting ratio was arbitrary. This produced the following weighted term set of 1339 terms, approximately 30 percent of the original dictionary, visualized in a word cloud in Figure 19. The final dictionary of terms is provided in Appendix 1.

Figure 19: Word cloud of the terms from a weighted average of the DistilBERT and SecureBERT Plus algorithms.

Understanding Existing Overlap Between Cyber and Resilience Terms

Hypotheses 3 and 4 seek to understand how well the existing cybersecurity documents incorporate aspects of resilience (H3) and rise above the technical and sociotechnical aspects (H4) of cybersecurity. Using the prior tf*idf and Latent Dirichlet Allocation analysis results, along with the lemmatized set of resilience definitions, the final dictionary can be analyzed for that overlap. This analysis of the dictionary against the resilience terms will provide a first look into how the answers to those hypotheses are trending.

The use of the prior analyses and lemmatized set of resilience definitions will provide a multi-faceted look at the dictionary. The lemmatized set of definitions provides a straightforward comparison and overlap, regardless of term importance or frequency. It answers the broader

question of whether the dictionary has any relationship to the resilience definitions, regardless of how the terms are used. The tf*idf and LDA analyses provide a focused overlap examination to determine how much overlap there is between the dictionary terms and the most important terms and topics from the resilience definitions.

Overall, the results of this analysis can be visualized as a four-set Venn diagram, indicating the amount of overlap between the sets. Figure 20, below, shows the results of this analysis.

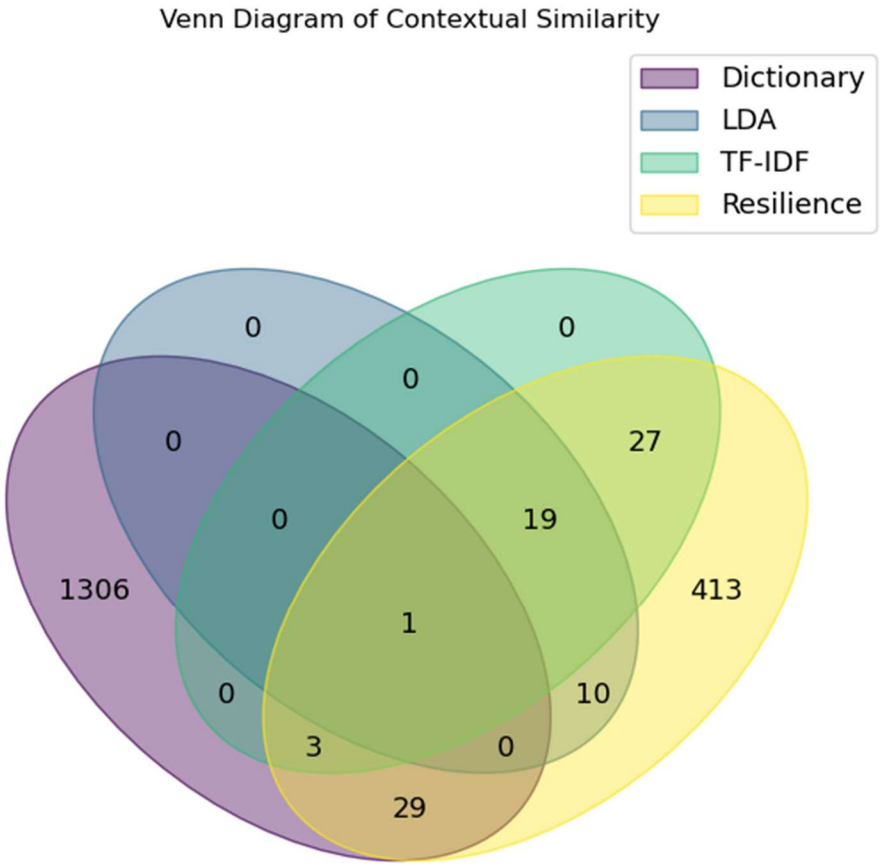

*Figure 20: A Venn diagram showing the overlap of terms between the combined cyber terms dictionary, the set of lemmatized resilience definitions, the tf*idf analysis results, and the Latent Dirichlet Allocation analysis results.*

As is readily apparent, there is little overlap between the resilience analyses and the combined set of cyber terms. The coverage ratio, a simple division of the total number of terms overlapping the dictionary divided by the dictionary, is found to be 2.46%. The following words overlapped with the dictionary:

- Tf*idf analysis (4): State, recovery, resource, and attack.
- LDA analysis (1): State.
- Lemmatized Resilience Definitions (33): Exposure, disruption, response, hazard, breach, target, learning, operation, knowledge, skill, state, group, goal, attribute, author, failure, attack, damage, degradation, capability, stage, vulnerability, problem, activity, asset, resource, stability, recovery, incident, relationship, survivability, compromise, and critical.

Given the research in Chapter 3 on the challenges of finding a resilience definition, it is possible that the resilience terms simply are not popular or relevant enough to have made it into the dictionary but may still be present in the broader corpus of 4,465 terms that were culled. The same analyses can be performed on that broader set to account for this possibility. The results are shown in Figure 21, below.

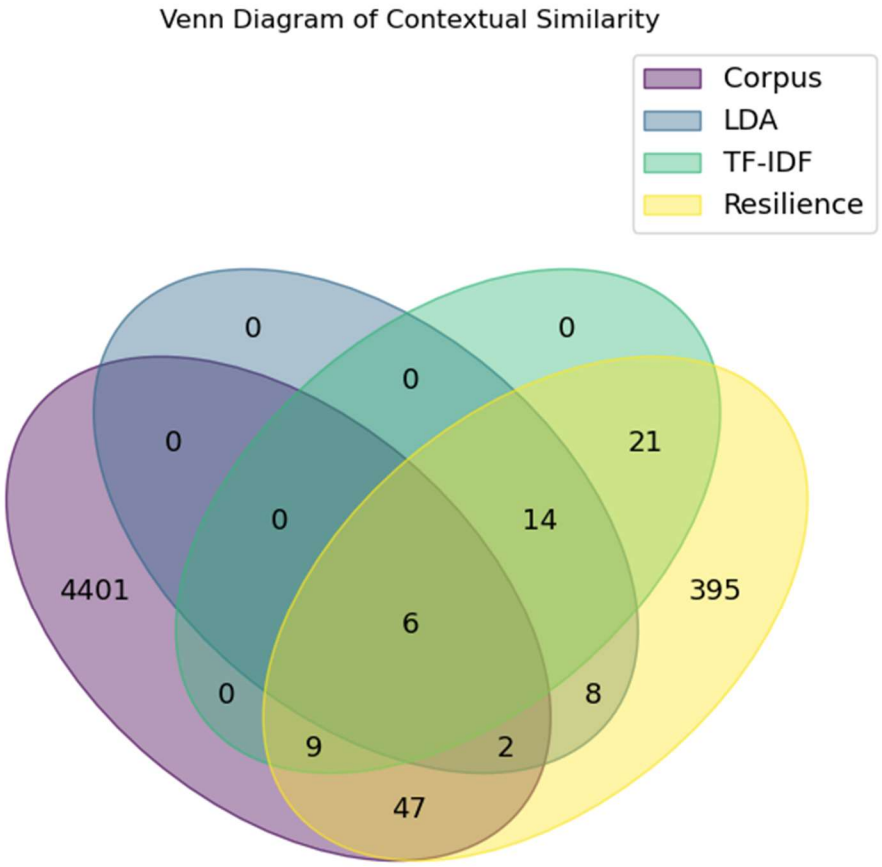

*Figure 21: A Venn diagram showing the overlap of terms between the total corpus of cyber terms, the set of lemmatized resilience definitions, the tf*idf analysis results, and the Latent Dirichlet Allocation analysis results.*

The coverage ratio for this broader set is lower, at 1.43%. The reduction in coverage ratio indicates that the overlapped terms are sufficiently relevant in the cyber term corpus to have passed selection into the final dictionary, and that there are not proportionally more terms present in the broader corpus that are unaccounted for. The overlapped terms are:

- Tf*idf analysis (15): State, recovery, resource, attack, environment, event, process, disturbance, cyber, function, risk, outcome, identity, individual, and adversity.
- LDA analysis (8): State, variable, function, disturbance, process, event, mitigate, and environment.

- Lemmatized Resilience Definitions (64): Exposure, environment, mitigate, challenge, disruption, response, event, hazard, device, risk, breach, identity, cybersecurity, target, bend, learning, operation, knowledge, process, skill, disturbance, cyber, parameter, network, core, outcome, state, group, goal, individual, attribute, failure, author, actor, specific, safety, attack, context, domain, aid, damage, adversity, degradation, threat, comparison, capability, stage, vulnerability, baseline, problem, identify, activity, asset, resource, stability, recovery, function, variable, incident, relationship, service, survivability, compromise, and critical.

As this analysis shows, both the curated dictionary and the broader corpus of cyber-related terms have minimal exposure to the resilience terms identified in the previous chapter. A portion of the resilience terms do appear in both the broader corpus and the final dictionary, but the lack of broader context for the terms in the dictionary prevents a more thorough analysis of co-occurrence and semantic relationships between the terms—are the resilience terms being used in ways that connote resilience as the interdisciplinary community has loosely defined it? That level of analysis will be saved for answering the central research question when the semantic richness in whole documents can be used.

The Analytical Scaffold and Term Classification by Time and Scale

To assess how existing guidance supports the development of cyber resilience, the concepts presented in this chapter can be used to create a conceptual scaffolding or framework from which to analyze existing guidance. At the meta level, resilience has components across both time and scale, as previously described in Chapter 4. These can be used to develop a two-

dimensional scale similar to Figure 11 in Chapter 3. The use of a two-dimensional scale will allow for both qualitative and quantitative analysis of the existing cyber guidance. Additionally, the use of a two-dimensional scale also provides the start of a useful mental model for practitioners to think through qualitatively, and potentially quantitatively, how to assess the level of resilience in an organization or develop a strategy to improve cyber resilience.

The two-dimensional scaffold assesses each term for its localized point of action with respect to time and scale. For example, a firewall operates at machine speed (time) and is a part of the technical system (scale), whereas a cyber risk management and governance process in an organization operates at human speeds (time), measured in days to weeks, and primarily at the organizational level (scale). The qualitative classifiers will be assigned numerical values solely for the purpose of quantitative analysis later, such as centroid calculations based on scatterplots of the terms and similar approaches. To account for the log-log relationship identified in Chapter 4, and to provide further separation for quantitative analysis, the numerical values assigned to each category of time and scale will be taken from the Fibonacci sequence rather than sequential integers. Table 6, below, shows the classifier assignments. Further details on what each classified generally covers will follow.

Table 6: Assignment of classifiers for time and scale attributes. Numerical assignment follows the Fibonacci sequence.

Numerical Assignment	Time	Scale
2	Immediate	Technical
3	Short	Sociotechnical
5	Mid	Organizational
8	Long	Community/Sector
13	Generational	National/International

For the time classifiers, the terms connote the following:

- Immediate: Seconds to minutes, machine speed, instantaneous, automated. May apply to terms like automated logging, real-time alerts, intrusion detection systems, firewalls, etc.
- Short: Hours to days, benchmarked off human capacity for response and adaptation. May apply to things like incident response, software patching, DevOps, security information and event management (SIEM) systems, etc.
- Mid: Weeks to months, the time it takes organizations to implement longer term changes. May apply to enterprise architecture updates, risk management, supply chain management, policy changes, threat hunting, etc.
- Long: Years. These are system shifts or adaptations that evolve over an extended period, such as implementation of new national or international regulations or standards. May apply to public laws, governance and risk compliance activities, executive orders, long-term strategy, etc.
- Generational: Multiple decades to centuries. Not likely to be seen for cyber resilience applications, but would apply to things like climate change, geopolitical shifts, evolution of technology and norms, ecological changes, etc.

For the scale classifiers, the terms generally apply as follows:

- Technical: The focus is on individual technical components or a technical system that do not require significant human interaction. May apply to terms like firewall, cryptographic keys, wide area network, hard drive, etc.
- Sociotechnical: The combination of technical systems that require human interaction to function properly. May apply to terms like identity and access management, access

controls, incident response, security information and event management systems, DevOps pipelines, software development, threat hunting, vulnerability management, etc.

- Organizational: The human side of the sociotechnical system, where primarily human-driven activities take place. This includes terms like risk management, supply chain management, governance, policy development, organizational strategy, organizational culture, etc.
- Community/Sector: How the organization fits into a broader ecosystem. The ecosystem has a generally identifiable set of norms, practices, and governance structure. This includes terms like sector-specific policies, critical infrastructure, ecosystems, industry standard development, interoperability, information sharing and analysis centers, etc.
- National/International: Focuses primarily on actions that operate on a national or international scale. May apply to terms like national resilience, executive orders, United Nations or other international body actions, such as the Convention on Climate Change (COP29) framework on reduction targets, etc. Language here is expected to be far less technical and more broadly applicable than for any other scale.

To assist in classification across both dimensions for each for the 1,339 terms, the above information was fed into the Claude 3.5 Sonnet algorithm to generate example classifiers for each scale to support follow on few-shot classification [132].[7] The example classifiers are used in with a Jaccard similarity calculation to determine overlap with terms to classify the terms by time, scale, and numerical assignments, and export the results to a file for review. The output of the classification script was manually reviewed to check for reasonableness of classifier

[7] Several algorithms were tried, including ChatGPT and DistilBERT (via a python script), but Claude 3.5 Sonnet proved more capable for this application of zero-shot application than the other algorithms.

assignment and its corresponding numerical assignment, and the classification was performed iteratively to adjust example language. Several small, final adjustments were made manually to arrive at a final classification and numerical assignment. The final dictionary, with classifiers but without numerical assignments, is provided in Appendix 1 for further review.

Curating Cyber and Resilience Guidance and Standards for Evaluation

The central research question (RQ1) seeks to determine whether a single strategy, standard, or other guidance document exists to provide organizations at multiple levels of society with support for improving cyber resilience. An expansive search was conducted to gather as many different cyber standards and other guidance documents as possible. 37 documents were gathered, including those from U.S and other international government agencies, standards organizations, and other sources, and from highly technical publications such as firmware resiliency to holistic or sector-specific standards, such as critical infrastructure or cyber supply chain risk management and cyber-resilient systems development:

- ISO/IEC 22301: Security and Resilience [135]
- ISO/IEC 27001: Information Security, Cybersecurity and Privacy Protection [136]
- ISO/IEC 27002: Information Security [137]
- ISO/IEC 31000: Risk Management Guidelines [138]
- NIST Cybersecurity Framework Version 2.0 [139]
- NIST SP 800-37: Risk Management Framework for Information Systems and Organizations [140]
- NIST SP 800-50: Building a Cybersecurity and Privacy Learning Program [141]

- NIST SP 800-53: Assessing Security and Privacy Controls in Information Systems and Organizations [31]
- NIST 800-61: Incident Response Recommendations and Considerations for Cybersecurity Risk Management [142]
- NIST 800-82: Guide to Operational Technology Security [143]
- NIST SP 800-160 Volume 1: Systems Security Engineering [144]
- NIST SP 800-160 Volume 2: Developing Cyber-Resilient Systems [28]
- NIST SP 800-161: Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations [145]
- NIST SP 800-193: Platform Firmware Resiliency Guidelines [146]
- NIST SP 800-207: Zero Trust Architecture [147]
- NIST SP 800-221A: Information and Communications Technology Risk Outcomes [148]
- NIST SP 1800-35: Implementing Zero Trust Architecture [149]
- National Cyber Security Centre Cyber Assessment Framework Version 3.2 (United Kingdom) [150]
- BSI Standard 100-1: Information Security Management Systems (Germany) [151]
- BSI Standard 200-1: Information Security Management Systems (Germany) [152]
- BSI Standard 200-2: IT-Grundschutz Methodology (Germany) [153]
- BSI Standard 200-3: Risk Analysis Based on IT-Grundschutz (Germany) [154]
- BSI IT-Grundschutz Compendium (Germany) [155]
- Center for Internet Security Critical Security Controls Version 8.1 [156]
- Department of Defense Cybersecurity Maturity Model Certification Version 2.0 [157]
- COBIT 5 Framework [158]

- Axelos ITIL 4: Framework for the Management of IT-Enabled Services [159]
- Department of Energy – Energy Sector Cybersecurity Framework Implementation Guidance [160]
- Federal Financial Institutions Examination Council Cybersecurity Assessment Tool [161]
- OCTAVE Criteria Version 2.0 [162]
- Executive Order 14028 on Strengthening and Promoting Innovation in the Nation's Cybersecurity [163]
- World Economic Forum Cyber Resilience Index [164]
- MITRE Cyber Resiliency Engineering Aid [165]
- MITRE Cyber Resiliency Framework and Cyber Survivability Attributes [166]
- Lockheed Martin Cyber Resiliency Level [167]
- Business Resilience Council Operational Resilience Framework Version 2.0 [168]
- ENISA Cybersecurity Culture Guidelines [169]

To check the scaffolding for adequacy and provide juxtaposition for the cyber frameworks, standards, and other guidance, an additional twelve sources were gathered from other disciplines such as climate change and disaster response, national resilience, community resilience, and social-ecological resilience, for a total of 49 documents. The search used iterative prompting with GPT-4 to identify standards or other resilience strategy documents that would not be forthcoming as easily through a traditional internet search [11]. This human-machine teaming demonstrated that more than half of the documents below have similar features to the other resilience documents and cybersecurity frameworks.

- National Resilience Strategy (United States) [170]

- National Resilience in a Changing Security Environment (Ukraine) [171]
- National Climate Resilience Framework (United States) [172]
- World Health Organization: Health Emergency and Disaster Risk Management Framework [173]
- Basel Committee on Banking Supervision: Principles for Operational Resilience [174]
- Resilience Alliance: Assessing Resilience in Social-Ecological Systems [175]
- Community Resilience System Initiative Steering Committee – Final Report [176]
- Asia-Pacific Economic Cooperation Disaster Risk Reduction Framework [177]
- North American Electric Reliability Corporation Reliability Standards for the Bulk Electric Systems of North America [178]
- Cybersecurity & Infrastructure Security Agency Infrastructure Resilience Planning Framework [179]
- Cybersecurity & Infrastructure Security Agency National Infrastructure Protection Plan [180]
- Organization for Economic Cooperation and Development Guidelines for Resilience Systems Analysis [181]

As described in Chapter 2, the documents were passed through a curation pipeline to clean the documents, remove stop words, and tokenize and lemmatize the text. Stop word removal and tokenization was accomplished with the `NLTK` (V3.8.1) package along with a custom stop word dictionary developed through an iterative approach during the resilience document pre-processing and further customized based on additional iterative analysis with this new document set [16], [22]. The tokenized definitions were lemmatized using the `spaCy`

(V3.7.5) package with a maximum character length set to 5,000,000 to account for the longest texts [18].

A Data-Centric Understanding of the Corpus of 37 Cyber Documents

The remainder of this chapter analyzes the documents using tf*idf, Latent Dirichlet Allocation, cosine similarity, and bigram analyses, on the full corpus, the 37 cyber texts, and the twelve non-cyber texts, to understand the language in the documents. Splitting the analyses between the two groups of text allows for greater understanding on how resilience terms appear in cyber and non-cyber, but resilience related, texts. Classification with the previously developed scaffolding and the analysis of those classification results will be performed in the next chapter.

Hypothesis 4 postulated that the cyber guidance, frameworks, standards, and other strategies primarily involve the technical and related sociotechnical aspects of cyber security, not the broader aspects of resilience as understood by the broader interdisciplinary community. The analyses in this section indicate that the cyber documents in this study focus on risk management as the central factor in achieving cybersecurity objectives. The analyses used for the upcoming analyses use the following python packages: `glob` (V0.1.5), `json` (V3.13), `pandas` (V2.1.4), `numpy` (V1.24.3), `matplotlib` (V3.8), `seaborn` (V0.12.2), `scikit-learn` (V1.4.2), `wordcloud` (V1.9.3), and `NLTK` (V3.8.1) [16], [21], [22], [71], [72], [125], [126], [127], [130], [131], [182]. LDA analyses were confined to ten topics with a limit of twenty iterations to converge on the set of topics. As discussed in Chapter 3, ten topics is a reasonable starting point. Heuristic and qualitative analysis was iteratively conducted on the LDA results to ensure adequate topic separation content coverage. Ten topics proved sufficient to describe each of the

small texts and the corpus as a whole without much overlap. The results from each analysis will be briefly described, followed by a more thorough and integrative analysis of each body of texts.

Time Frequency-Inverse Document Frequency Analysis

Each of the documents was analyzed through the tf*idf processing pipeline, and the results at the document and corpus level were collected. At the level of an individual document, the statistical descriptors in Table 7 below indicate that most terms did not have meaningful significance within the body of the text. The mean score was close to zero, though the average (mean) maximum score for a document was 0.863. Many of these documents are short compared to broader machine learning data sets, so lower scores are expected since the terms simply cannot repeat often enough to emerge with a higher score as they would be able to in a larger data set. To understand the aggregate behavior across the 37 cyber texts, though, the scores for terms were summed and the top 25 results printed. The results of those analyses are shown in Table 8 and Figure 22 below.

*Table 7: Statistical descriptors for the histogram of tf*idf scores from the cyber text corpus*

Descriptor	Value
Mean	0.000957
Standard Deviation	0.007864
Minimum	0.0
Maximum	0.863318

*Table 8: Top 25 terms by combined tf*idf score from the cyber text corpus*

Term	Td-idf Score
Risk	7.675829
Security	6.618338
Information	5.538127
Organization	4.971197
Management	3.874349
Process	3.307799
Include	2.985525
Service	2.928825
Control	2.747019
Access	2.297922
Enterprise	2.270666
Requirement	2.212232
Cybersecurity	2.139381
Provide	2.137164
Cyber	2.116034
Organization	2.012935
Business	1.928364
Device	1.900593
Resource	1.848839
Datum	1.826995
Identify	1.813171

Network	1.812263
Level	1.803199
Need	1.697559
User	1.687856

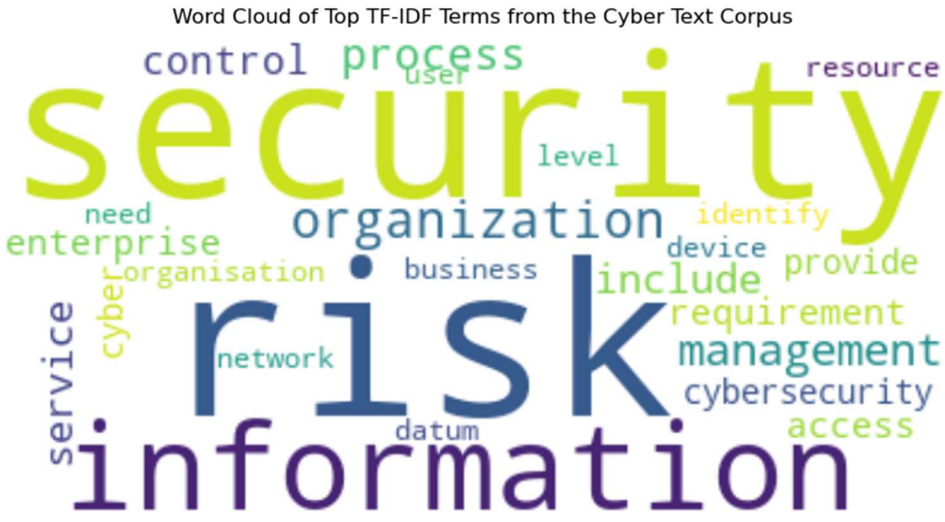

*Figure 22: Word cloud of the top 25 tf*idf terms from the corpus of 37 cyber texts.*

As the word cloud and top terms indicate, the largest themes across the body of 37 cyber texts are in risk and information management within an organization. While the efficacy of risk management procedures will not be explored in this research, the prior literature in this area calls its effectiveness into question [183]. Tellingly, the term resilien* does not appear in the top 25 terms, despite several of the texts identifying resilience as a stated objective of the guidance.

Latent Dirichlet Allocation Analysis

The Latent Dirichlet Allocation analysis results for the 37 cyber texts cover the major themes of the current best practices in organizational cyber security. The ten topics are provided in Table 9. As an unofficial or heuristic evaluation, the topics generally align with the Certified Information System Security Practitioner (CISSP) and the Security+ topical domains and body of knowledge, as should be expected with certifications demonstrating a solid fundamental understanding of the core concepts and best practices in organizational cybersecurity.

Table 9: Topics and interpretations from a Latent Dirichlet Allocation analysis of the cyber text corpus

Topic	Terms in the Topic	Overall Theme
1	Carry, cyber, transfer, experience, introduce, proper, arise, organizational	Organizational cybersecurity
2	Security, information, requirement, support, human, organization, consideration, zone, offer, frequently	Security requirements and support
3	Authority, treatment, role, risk, public, position, asset, review, select, package	Asset management and authorization
4	Stage, risk, construct, team, market, deployment, language, format, significant, overall	Project management and implementation
5	Member, fundamental, detect, ability, enhance, patch, locate, perspective, authentication, reference	Security enhancement and authentication

6	Security, time, consideration, management, information, responsibilities, strategy, flow, room, significantly	Security strategy
7	Inventory, problem, personnel, interact, provision, future, skill, integrity, decide, connection	Personnel and system integrity
8	Risk, security, information, organization, management, process, service, include, control, enterprise	Enterprise security controls
9	Similar, management, expectation, trade, intrusion, profile, asset, useful, conflict, large	Asset and intrusion management
10	Statement, artifact, staff, capture, requirement, encryption, accord, organize, hand, coordinate	Staff and security implementation

Cosine Similarity Analysis

Cosine similarity is a measure of relatedness between two texts. The similarity values can fall between zero and one, with zero representing no overlap or similarity between the documents, and one representing a high degree of similarity. In the analysis of the 37 cyber texts, the majority of the cosine similarity scores, shown in the histogram in Figure 23 below fall between 0.1 and 0.4, indicating that the documents do not possess a high degree of similarity and

maintain their distinctiveness. The low similarity scores suggest that the corpus of 37 cyber texts will provide a broader perspective on cyber resilience than if the scores had a higher degree of similarity.

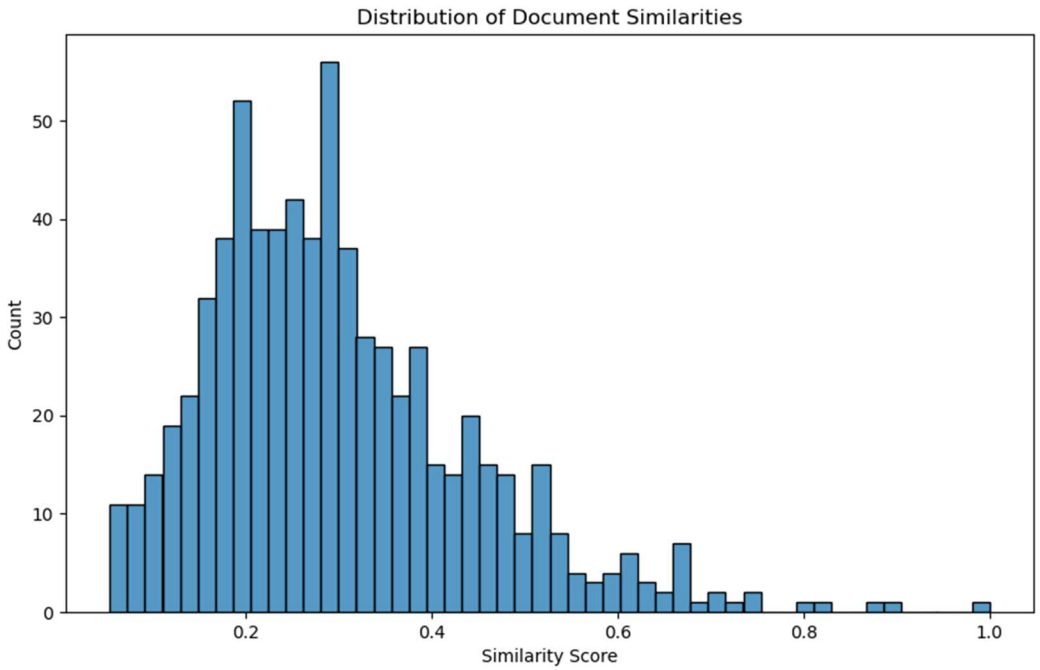

Figure 23: Histogram of cosine similarities for 37 cyber texts.

The pair index, Figure 24, shows the rank ordering of similarity scores, revealing the distribution shape and any outliers. While the histogram shows how many pairs have each similarity score, the pair index plot shows which pairs represent the highest and lowest similarities and how gradually or sharply the similarities change across all document pairs. This helps identify clusters of similarly-related documents and spot unusual document relationships that might be hidden in histogram bins. The "pair index" in the scatter plot represents the index of the document pairs when the similarity values are flattened into a single list. The positions of

the similarity values and shape of the pair index provide an alternative visual interpretation of the information in this histogram, particularly in highlighting how rapidly the paired similarity scores change across the corpus. In this case, as the histogram in Figure 23 indicates, most of the pairwise similarity scores indicate low similarity.

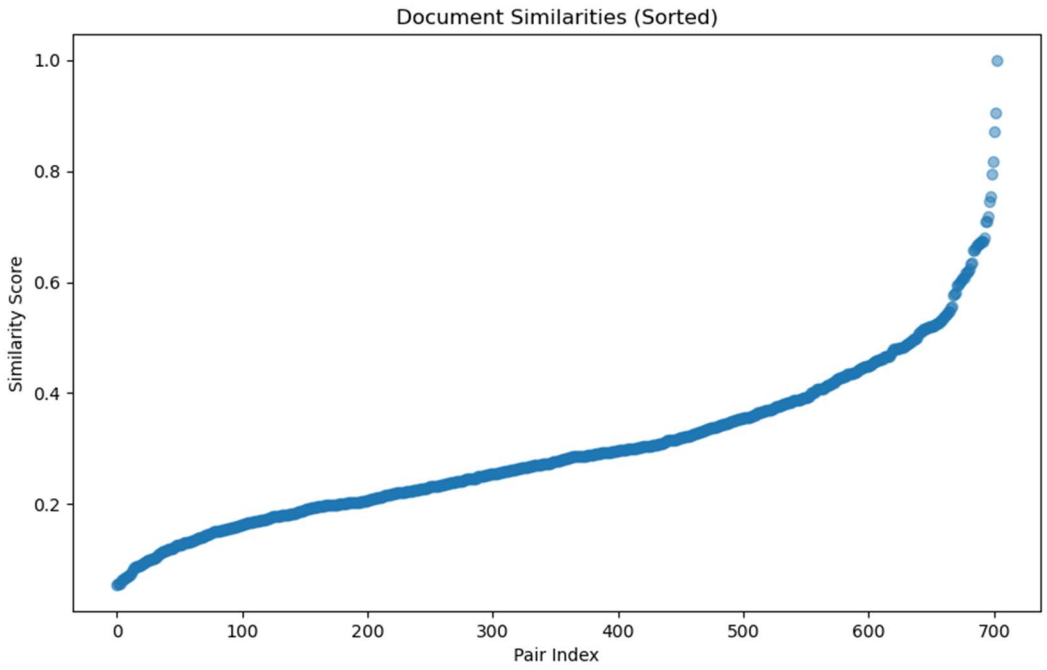

Figure 24: Pair index of sorted cosine similarities showing overall similarity trends in the corpus of 37 cyber texts.

Frequency of Word Pairings: Bigram Analysis

Finally, a bigram analysis of the texts indicates how often a particular word appears with another word and provides greater insight into how a particular word tends to be used in the texts. This analysis was performed on the 37 texts for the top bigrams containing the words "resilience," "risk," and "security," since the latter two words were dominant in the tf*idf analysis from Table 8. For the bigrams containing "resilience," the analysis is performed for both

the top bigrams containing “resilience” and the top bigrams in the documents with a higher tf*idf score for resilience.

Table 10: Bigram analysis for "resilience" of 37 cyber texts, for bigrams containing "resilience" and document frequency for texts containing "resilience"

Term: Resilience	Frequency	Term: Resilience	Document Frequency
Cyber resilience	74	Risk management	33
Operational resilience	41	Risk assessment	32
System resilience	20	Information system	31
Security resilience	19	Information security	30
Organizational resilience	12	Internal external	28
Resilience engineering	12	Information technology	28
Resilience principle	9	Confidentiality integrity	27
Resilience safety	8	Security control	27
Culture resilience	8	Management system	26
Resilience survivability	7	Management process	26

Table 11: Top 10 bigrams containing "risk" or "security" in the 37 cyber texts and their frequency

Term: Risk	Frequency	Term: Security	Frequency
Risk management	1573	Information security	2017
Risk assessment	624	Security privacy	1346
Cybersecurity risk	478	System security	782
Risk analysis	339	Security aspect	730
Supply risk	337	Security requirement	693

Privacy risk	310	Security policy	566
Chain risk	299	Security concept	454
Security risk	267	Security safeguard	412
Risk response	221	Security incident	412
Risk supply	203	Security objective	351

Table 10 and Table 11 show a pronounced contrast in how the same set of texts treat resilience with request to the common cybersecurity terms “risk” and “security.” The frequency with which the documents discuss more traditional terms compared to resilience terms indicates, despite the titles or stated objectives, these documents are not likely to address resilience concepts at depth, or in a manner commensurate with the broader interdisciplinary community. As Table 10 shows, the majority of the cyber texts discuss resilience in some capacity, but from the documents that do discuss resilience, the top bigrams from those documents are traditional cybersecurity terminology.

A Data-Centric Understanding of the Non-Cyber Document Corpus

The same analysis was performed on the twelve non-cyber documents to get a comparative understanding against the 37 cyber documents. The tf*idf, Latent Dirichlet Allocation, and bigram analyses indicate, as expected, a stronger focus on the concepts of resilience. These documents were selected as a pseudo-control group since they were explicitly resilience focused and originated from non-cyber perspectives. Of note, the results from these analyses suggest that this corpus of texts also falls somewhat short of the interdisciplinary

definition of resilience and its related concepts, though it is clear that they are more closely aligned than the cyber texts.

Time Frequency-Inverse Document Frequency Analysis

At the level of an individual document, the statistical descriptors of the non-cyber texts, in Table 12 indicate that most terms did not have meaningful significance within the body of the text. The mean score was close to zero, though the average (mean) maximum score for a document was 0.722. Many of these documents are short compared to broader machine learning data sets, so lower scores are expected since the terms simply cannot repeat often enough to emerge with a higher score as they would be able to in a larger data set. The top 25 terms from the tf*idf analysis are shown in Table 13, below, and visualized as a word cloud in Figure 25.

*Table 12: Statistical descriptors for the histogram of tf*idf scores from the cyber text corpus*

Descriptor	Value
Mean	0.001655
Standard Deviation	0.009891
Minimum	0.0
Maximum	0.722077

*Table 13: Top 25 terms by combined tf*idf score from the cyber text corpus*

Term	Td-idf Score
Resilience	2.693974
Community	2.347767

Risk	1.961130
Infrastructure	1.334006
Include	1.224391
Critical	1.140232
Provide	0.947927
National	0.926513
Plan	0.892367
Change	0.888363
Health	0.817340
Climate	0.813726
Disaster	0.807247
Security	0.805513
Management	0.793606
Support	0.742548
Identify	0.727888
Impact	0.714154
Resource	0.711756
Sector	0.680041
Information	0.679013
Process	0.652294
Need	0.593144
Ensure	0.592306
Emergency	0.590405

*Figure 25: Word cloud of the top 25 tf*idf terms from the corpus of twelve non-cyber texts.*

The tf*idf and word cloud analysis results indicate a much different corpus of texts than for the 37 cyber texts. The twelve non-cyber texts place a greater emphasis on change and factors beyond the organization's control than for the 37 cyber texts.

Latent Dirichlet Allocation Analysis

The Latent Dirichlet Allocation analysis results in Table 14, contrary to the results in Table 9 for the 37 cyber texts, implicitly emphasize adaptation, change, and cross-scale effects. The set of topics generally addresses the time and scale aspects of resilience identified in the previous chapter. For scale, the topics address organizational, community, national, and global

aspects of resilience. For time, the topics address near-term disaster recovery to longer term recovery and adaptation.

Table 14: Topics and interpretations from a Latent Dirichlet Allocation analysis of the non-cyber text corpus

Topic	Terms in the Topic	Overall Theme
1	Channel, setting, trip, follow, leadership, receive, pc, circumstance	Leadership and operations
2	Owner, group, infrastructure, implementation, contract, funding, territorial, retire, traditional	Infrastructure and funding
3	Baseline, second, recovery, type, process, attribute, significant, schedule, place, stage	Recovery process
4	Strong, threat, resilience, coordinated, modify, distance, line, infrastructure, specific, grant	Infrastructure security
5	Natural, change, hazard, resilience, ability, bank, early, experience, community, mean	Community resilience
6	Resilience, community, risk, infrastructure, critical, include, national, plan, security, provide	National security planning

7	Legal, help, ramp, language, recover, prevent, staff, determine, trust	Legal recovery
8	Characteristic, fund, flexibility, resilience, shock, exist, climate, guidance, safe, minute	Climate and resilience
9	Specification, focus, transition, leave, miss, reactive, utilize, context, infrastructure, list	Infrastructure transitions
10	Health, include, emergency, subsequent, stress, challenge, people, attachment, excursion, retire	Emergency and health response

Cosine Similarity Analysis

The cosine similarity analysis shows that the twelve non-cyber texts have equivalent similarity metrics with the 37 cyber texts. The bulk of the documents had similarity scores between 0.2 and 0.4, with all documents falling between 0.1 and 0.7, as showing in the histogram in Figure 26. The pair index results for the twelve non-cyber texts, shown in Figure 27, indicate that the corpus of texts shares low similarity scores. The generally linear progression of the pair index shows no clustering of documents that would suggest greater similarity between a subset.

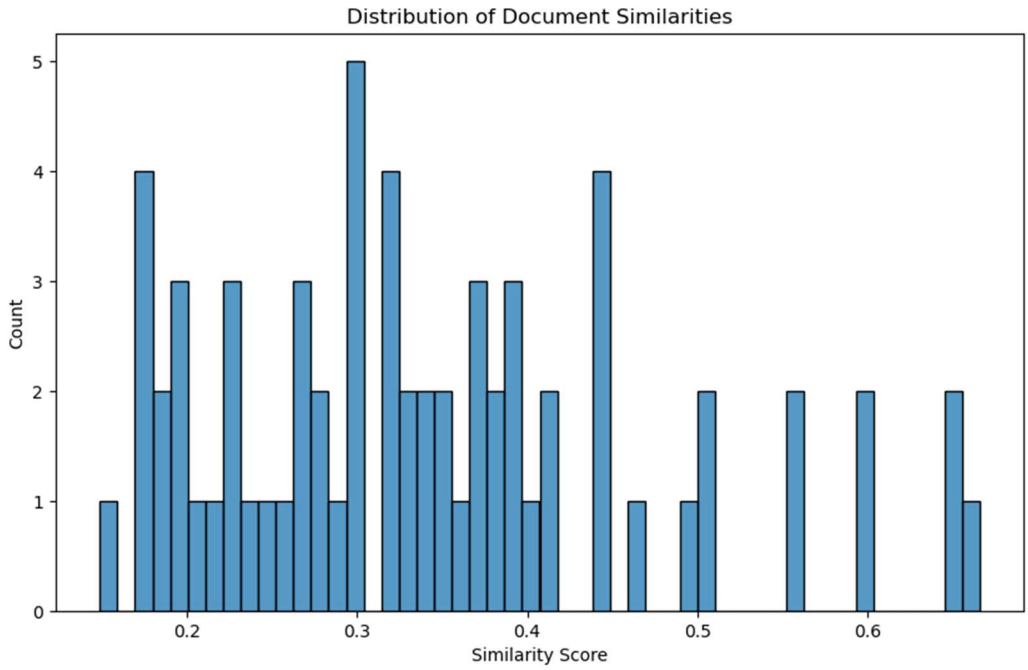

Figure 26: Histogram of cosine similarities for twelve non-cyber texts.

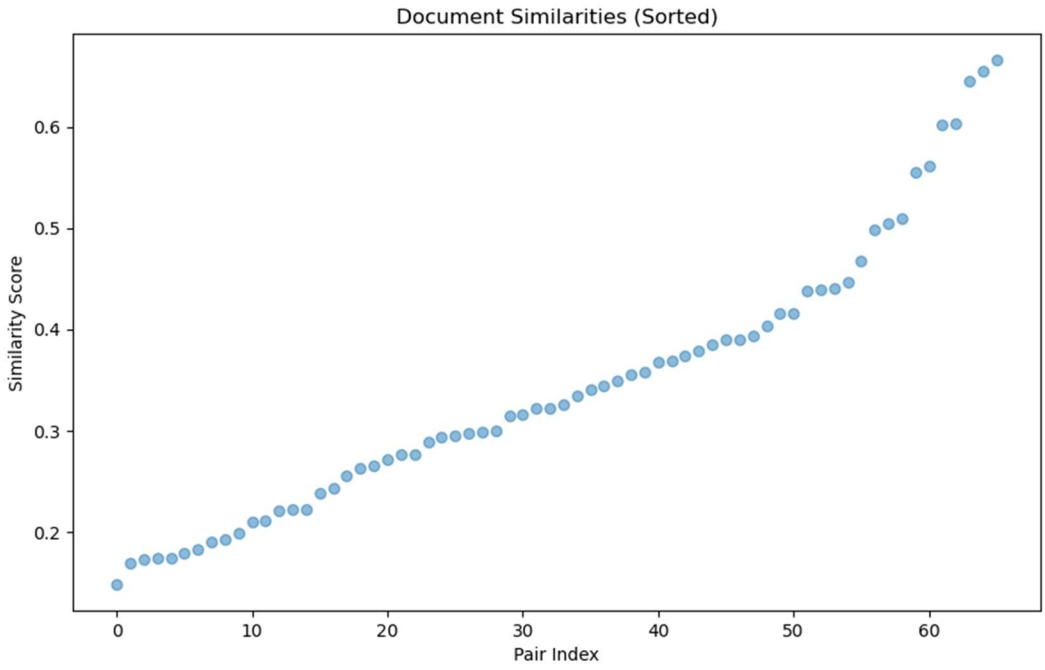

Figure 27: Pair index of sorted cosine similarities showing overall similarity trends in the corpus of twelve non-cyber texts.

Frequency of Word Pairings: Bigram Analysis

The same bigram analysis was repeated for the twelve non-cyber texts, with the results shown in Table 15 and Table 16, below. The bigrams strike a different tone than for the 37 cyber texts, with an emphasis on both traditional risk management and broader resilience concepts. The terms also appear in greater numbers, as expected for resilience-focused documents. Notably, the bigram "Security resilience" appeared in the list of top bigrams for the term "security," which indicates a likely connection between security and resilience outcomes in the documents.

Table 15: Bigram analysis for "resilience" of twelve non-cyber texts, for bigrams containing "resilience" and document frequency for texts containing "resilience"

Term: Resilience	Frequency	Term: Resilience	Document Frequency
National resilience	678	Risk assessment	11
Security resilience	360	Public private	10
Resilience ensure	306	Risk management	10
Community resilience	187	Good practice	10
Resilience system	121	System include	10
Resilience concept	83	Private sector	9
Climate resilience	74	Community region	9
System resilience	62	Natural hazard	9
Enhance resilience	54	Supply chain	9
Ensure resilience	53	Risk reduction	9

Table 16: Top 10 bigrams containing "risk" or "security" in the twelve non-cyber texts and their frequency

Term: Risk	Frequency	Term: Security	Frequency
Risk assessment	276	National security	511
Risk management	171	Security resilience	360
Risk threat	147	Cyber security	206
Disaster risk	63	Security plan	133
Assess risk	58	Security environment	121
Security risk	54	Infrastructure security	119
Reduce risk	50	Security defense	90
Manage risk	45	Physical security	78
Risk reduction	42	Security policy	61
National risk	41	Security ensure	55

A Data-Centric Understanding of the Entire Document Corpus

Intuitively, the sum of the two sets of text documents, the 37 cyber texts and the twelve non-cyber texts, should be the sum of their group results to this point, but since several of the methods, in particular Latent Dirichlet Allocation and cosine similarities, require the full corpus of text, the results will be slightly different. This penultimate section repeats the same analyses from the two separate blocks of texts. The final section of this chapter will make integrative observations about the development of the texts and the heuristic-based results on how well the 37 cyber texts will address some or all aspects of cyber resilience.

Time Frequency-Inverse Document Frequency Analysis

At the level of an individual document, the statistical descriptors of the non-cyber texts, in Table 17, below, indicate that most terms did not have meaningful significance within the body of the text. The mean score was close to zero again, though the average (mean) maximum score for a document was 0.858—closer to the maximum for the 37 cyber texts. The top 25 terms from the tf*idf analysis are shown in Table 18, below, and visualized as a word cloud in Figure 28.

*Table 17: Statistical descriptors for the histogram of tf*idf scores from the cyber text corpus*

Descriptor	Value
Mean	0.007885
Standard Deviation	0.007094
Minimum	0.0
Maximum	0.858698

*Table 18: Top 25 terms by combined tf*idf score from the cyber text corpus*

Term	Td-idf Score
Risk	9.193893
Security	7.279186
Information	6.017220
Organization	5.132970
Management	4.481233
Include	3.977743
Resilience	3.831211

Process	3.815730
Service	3.352385
Control	2.960307
Provide	2.913416
Community	2.519514
Cyber	2.494043
Access	2.481754
Requirement	2.460009
Enterprise	2.425171
Identify	2.423940
Resource	4.413110
Cybersecurity	2.351744
Critical	2.298630
Business	2.296908
Level	2.265422
Support	2.201590
Plan	2.167644
Change	2.155858

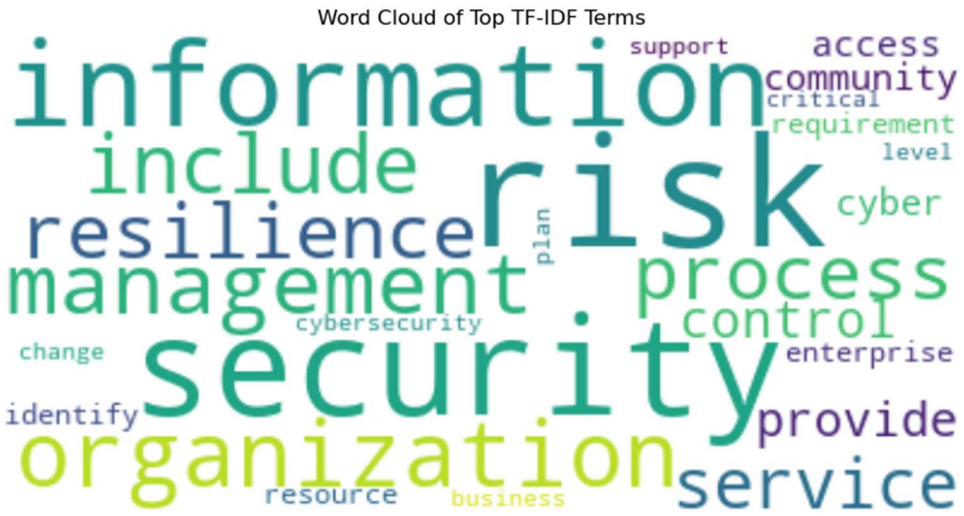

*Figure 28: Word cloud of the top 25 tf*idf terms from the corpus of 49 texts.*

While the word cloud remains generally the same compared to the word cloud and tf*idf results from the 37 cyber texts, in Figure 22 and Table 8, the inclusion of resilience, business, community, and change show the influence of the non-cyber texts. The broad corpus of documents seems to align with cybersecurity best practices and not be significantly influenced by the interdisciplinary resilience research.

Latent Dirichlet Allocation Analysis

For the full body of texts, the Latent Dirichlet Allocation analysis strikes a balance between traditional cybersecurity subject matter but does exhibit several aspects of resilience in the topic list. In particular, Table 19 highlights topics related to investment and transition and community resilience, and topic 5 on regulation could be defined as both cyber- and resilience-related.

Table 19: Topics and interpretations from a Latent Dirichlet Allocation analysis of the cyber text corpus

Topic	Terms in the Topic	Overall Theme
1	Center, sector, critical, fail, information, key, resilience, transfer, page, assess	Critical information assessment
2	Order, management, importance, zone, hardware, container, resource, follow, critical, supply	Resource management
3	Automate, room, reason, treatment, previous, community, security, perform, assistance, society	Community security
4	Stakeholder, continuity, minimize, category, technical, determination, lesson, furthermore, simple, path	Technical planning
5	Module, global, direction, ability, evaluate, place, major, posture, regulation, increase	Regulation
6	Risk, time, container, simulation, restriction, establish, enterprise, fulfill, rule, technical	Enterprise risk management
7	Lack, properly, investment, good, real, post, social, introduction, deliver, transition	Investments and transitions

8	Resilience, community, climate, central, focus, external, disturbance, emergency, retain, measurement	Community resilience
9	Risk, security, information, organization, management, include, process, resilience, service, control	Organizational security
10	Resource, assessment, station, stage, category, environmental, http, accordance, participate, cover	Environmental assessment

Cosine Similarity Analysis

The cosine similarity analysis does not reveal anything unexpected compared to the previous analyses of the individual sets of texts. The low similarity scores, with most falling between 0.1 and 0.4 (shown in Figure 29), indicate that each of the documents has a high degree of independence from the others. There is minimal redundancy between the different texts, even when adding the twelve non-cyber documents into the mix. Likewise, the pair index retains the same basic shape as the prior analyses showed (in Figure 24 and Figure 27). The lack of clustering in the pair index curve further reinforces the lack of independence between the documents, and the shallow slope follows the same shape as the individual results.

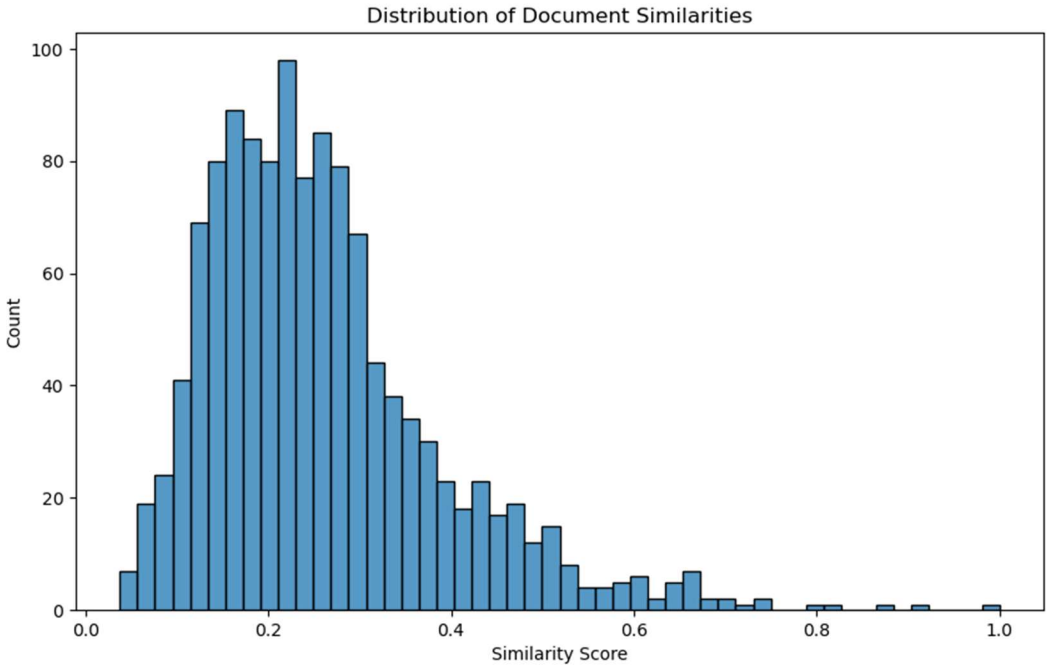

Figure 29: Histogram of cosine similarities for the 49-text corpus.

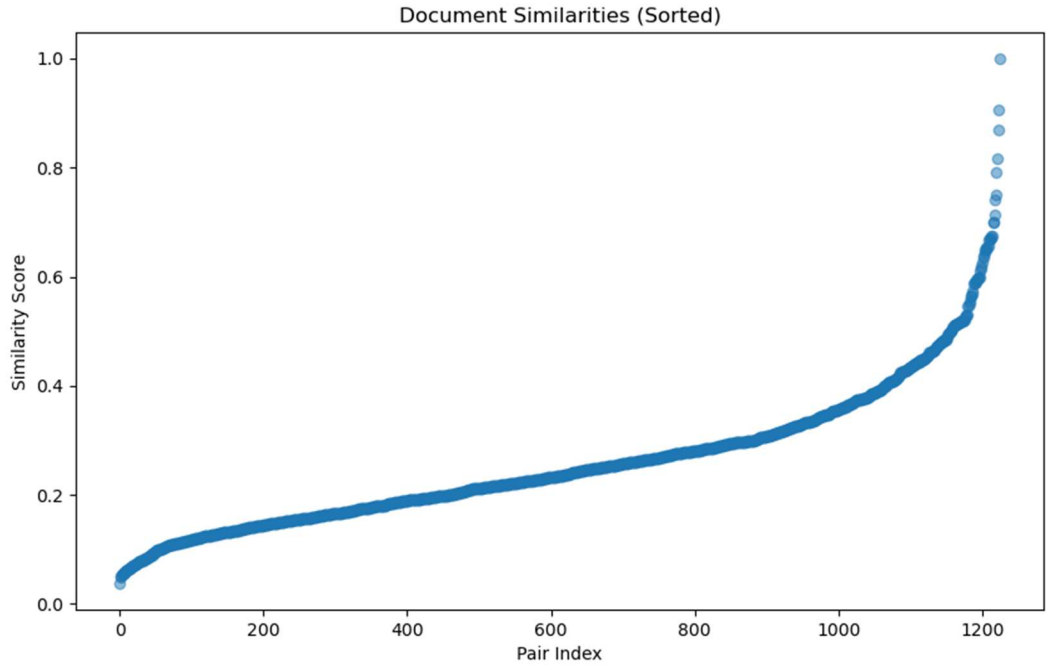

Figure 30: Pair index of sorted cosine similarities showing overall similarity trends in the corpus of 49 texts.

Frequency of Word Pairings: Bigram Analysis

The bigram analyses of the full body of documents shown in Table 20 and Table 21 skew heavily toward the traditional cybersecurity terminology. For documents that did mention resilience, the top ten bigram pairs in those documents are associated with core cybersecurity disciplines and practices. The addition of resilience-focused, non-cyber documents did assist in ensuring that cyber texts had the right pairings measured in them. Of note, the frequency with which major cybersecurity terms appeared indicates a strong tie between the cybersecurity texts, potentially diminishing aspects of resilience, which did not make the top ten terms for risk or security.

Table 20: Bigram analysis for "resilience" of all 49 texts, for bigrams containing "resilience" and document frequency for texts containing "resilience"

Term: Resilience	Frequency	Term: Resilience	Document Frequency
National resilience	678	Risk assessment	43
Security resilience	379	Risk management	43
Resilience ensure	307	Information system	39
Community resilience	187	Good practice	36
Resilience system	125	Supply chain	35
Resilience concept	85	Internal external	33
System resilience	82	Information technology	33
Operational resilience	81	Information security	33
Cyber resilience	74	Management system	32
Climate resilience	74	Management process	32

Table 21: Top 10 bigrams containing "risk" or "security" in all 49 texts and their frequency

Term: Risk	Frequency	Term: Security	Document Frequency
Risk management	1744	Information security	2022
Risk assessment	900	Security privacy	1347
Cybersecurity risk	480	System security	788
Risk analysis	363	Security aspect	730
Supply risk	337	Security requirement	702
Security risk	321	Security policy	627
Privacy risk	310	National security	597
Chain risk	302	Security concept	461
Risk response	222	Security safeguard	446
Risk supply	203	Security incident	414

Brief Assessment of the Data-Driven Analysis of the Corpus

This chapter began with a data-driven approach to building a dictionary of cyber and resilience terminology and finished with a holistic analysis of 49 cyber or resilience documents to gain insight into how well the words the interdisciplinary community uses with respect to resilience factor into cybersecurity guidance. The corpus of 37 cyber texts analyzed using tf*idf, word cloud analysis, Latent Dirichlet Allocation analysis, cosine similarity analysis, and bigram analyses reveal a body of cyber guidance that largely conforms to existing cybersecurity best practices.

With respect to the time and scale aspects of resilience developed earlier in this chapter, the corpus of 37 cyber texts qualitatively trends towards the machine and organizational levels

for time and scale, respectively. Analysis using the dictionary and classification framework will be performed in the next chapter. However, the semantic analyses performed here indicate that few to no documents discuss cyber resilience across time and scale in a way that would correspond with the interdisciplinary research into resilience. This leads to the tentative affirmation of Hypothesis 4 in that the documents focus primarily on technical and sociotechnical cybersecurity and lack guidance on most or all aspects of resilience.

CHAPTER 5: CLASSIFICATION RESULTS AND ANALYSIS

To answer the central research question—whether an existing strategy or other document exists to provide organizations at multiple levels of society with enough support for improving cyber resilience—Chapters 3 and 4 developed the interdisciplinary concept of resilience, a dictionary of cyber and resilience terminology, and a classification framework from which to evaluate 37 cyber-oriented texts against that question. The use of twelve non-cyber, resilience texts provides a pseudo-control group for comparison, since those documents intuitively should have higher resilience content than the cyber texts. The analysis in the previous chapter revealed that the cyber texts potentially validated research Hypothesis 4 since the semantic analyses indicated that, even though resilience was mentioned, the content tended to focus on the technical and sociotechnical aspects of cybersecurity only. To confirm that finding, this chapter will evaluate the texts individually, as opposed to collectively as a corpus, against the cyber resilience dictionary and the classification framework.

Classifying the Documents for Topics Associated with Time and Scale

The classification framework developed in the previous chapter evaluates each text on the basis of how well the terminology in the text corresponds to the descriptors across both the time and scale dimensions of resilience. The framework uses Fibonacci numbering to connote the non-linearity of each scale. As a recap, the time classifiers are coded as:

- 2: Immediate – seconds to minutes; machine speed
- 3: Short – hours to days, and the benchmark for human capacity
- 5: Mid = weeks to months

- 8: Long – an extended period, on the order of months, but more likely years
- 13: Generational – multiple decades to centuries

For the scale classifiers:

- 2: Technical – individual technical components with little to no human interaction
- 3: Sociotechnical – the combination of technical systems that require continual human interaction
- 5: Organizational – the human side of sociotechnical, where human-driven activities take place
- 8: Community/Sector – above the organization and how it fits into a broader ecosystem
- 13: National/International – actions at the national or international scale

Appendix 1 contains the final dictionary of terms with their time and scale classifiers.

The 49 texts were processed through a script that created a classification analyzer algorithm. The Python script used the following packages to implement: `os` (v3.13), `json` (v3.13), `matplotlib` (v3.8), `Seaborn` (v0.12.2), `pandas` (v2.1.4), and `NumPy` (v1.24.3) [125-130]. The texts were processed individually, and the mean and standard deviation of time and scale values stored as a `json` file for follow-on analysis. Each text, and the corpus of 49 texts, were visualized using a contour plot. The initial, combined plot of the results is shown below in Figure 31. The word clouds were generated with a separate script implementing a tf*idf analysis and using the above packages along with `scikit-learn` (v1.4.2) and `WordCloud` (v1.9.3) [71], [131].

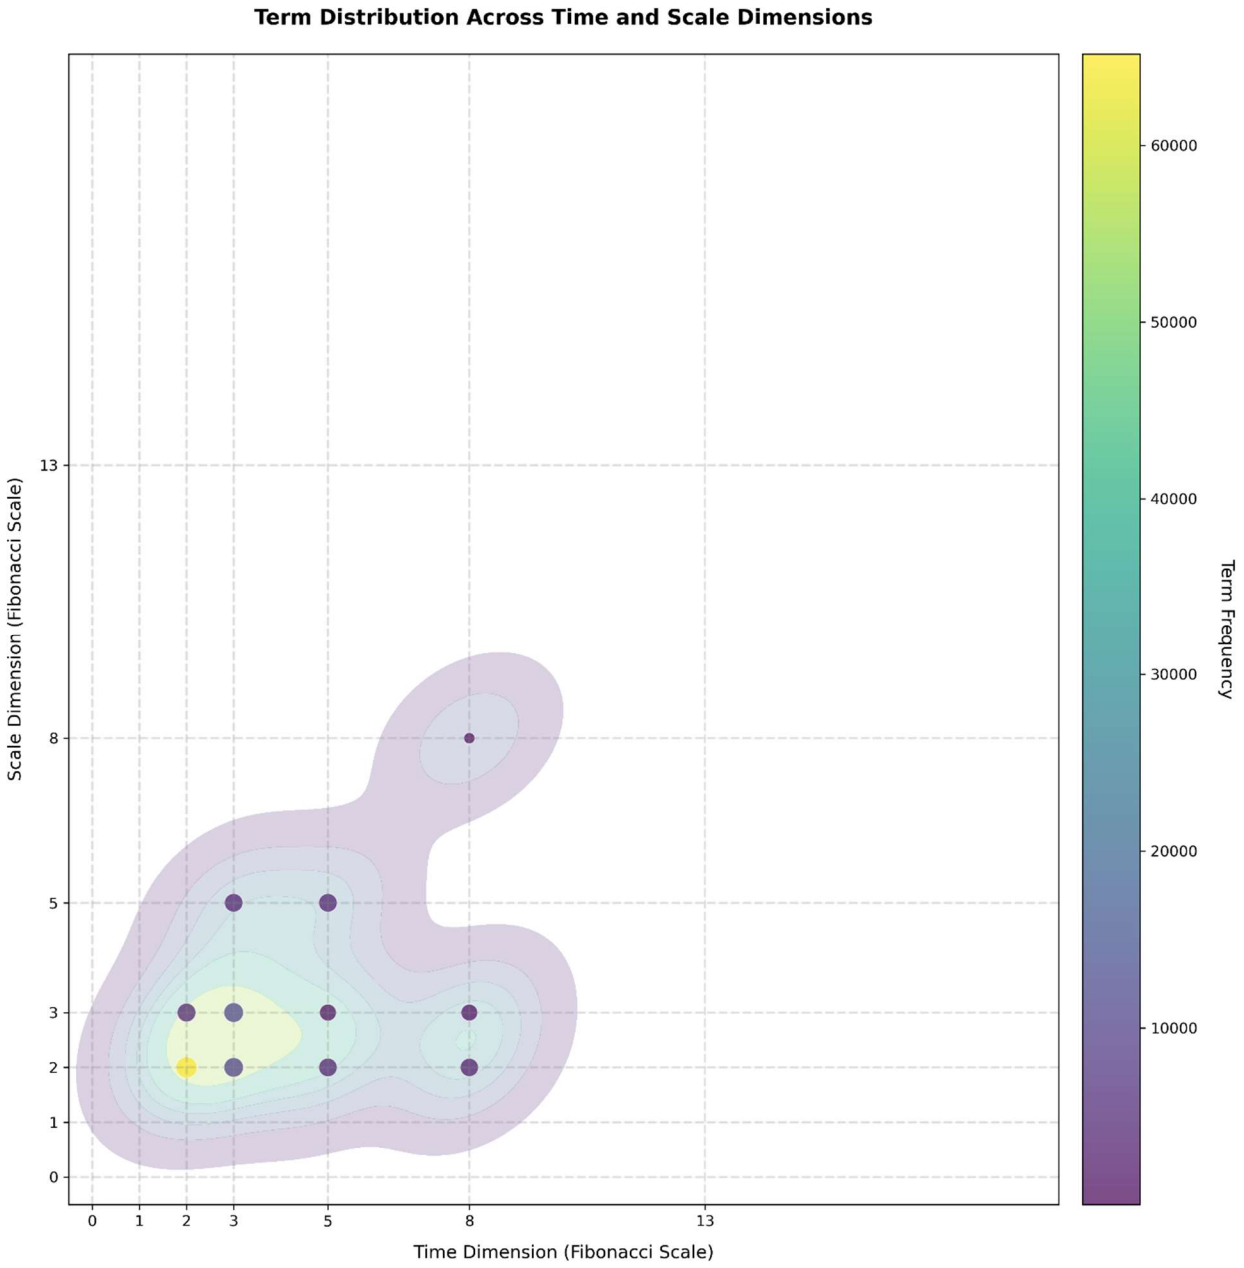

Figure 31: Contour plot with Fibonacci scaling representing the classification of 49 texts for aspects of resilience across time and scale.

As predicted in Chapter 2 with the motivating research question and Hypothesis 1, none of the existing body of documents, to include the twelve non-cyber texts used as a pseudo-control group, contain topics associated with resilience at all levels of time and scale. It is unreasonable to assume that specific cybersecurity guidance would contain discrete topics

associated with long-term or greater effects at the community/sector level or above. That should be the realm of national cyber resilience strategies. The resilience strategies for both the United States and Ukraine were in the twelve non-cyber texts, and the lack of coverage on the combined contour plot indicates that, at least according to the analytical procedures used in this research, even these strategies lack elements across time and scale related to cyber resilience. Figure 32 and Figure 33, below, show the contour plots for the cyber and non-cyber texts for comparison. Note the contouring scale differences between the two plots, though the same general distribution and dispersion of the features from classification remains visually consistent between the two bodies of texts.

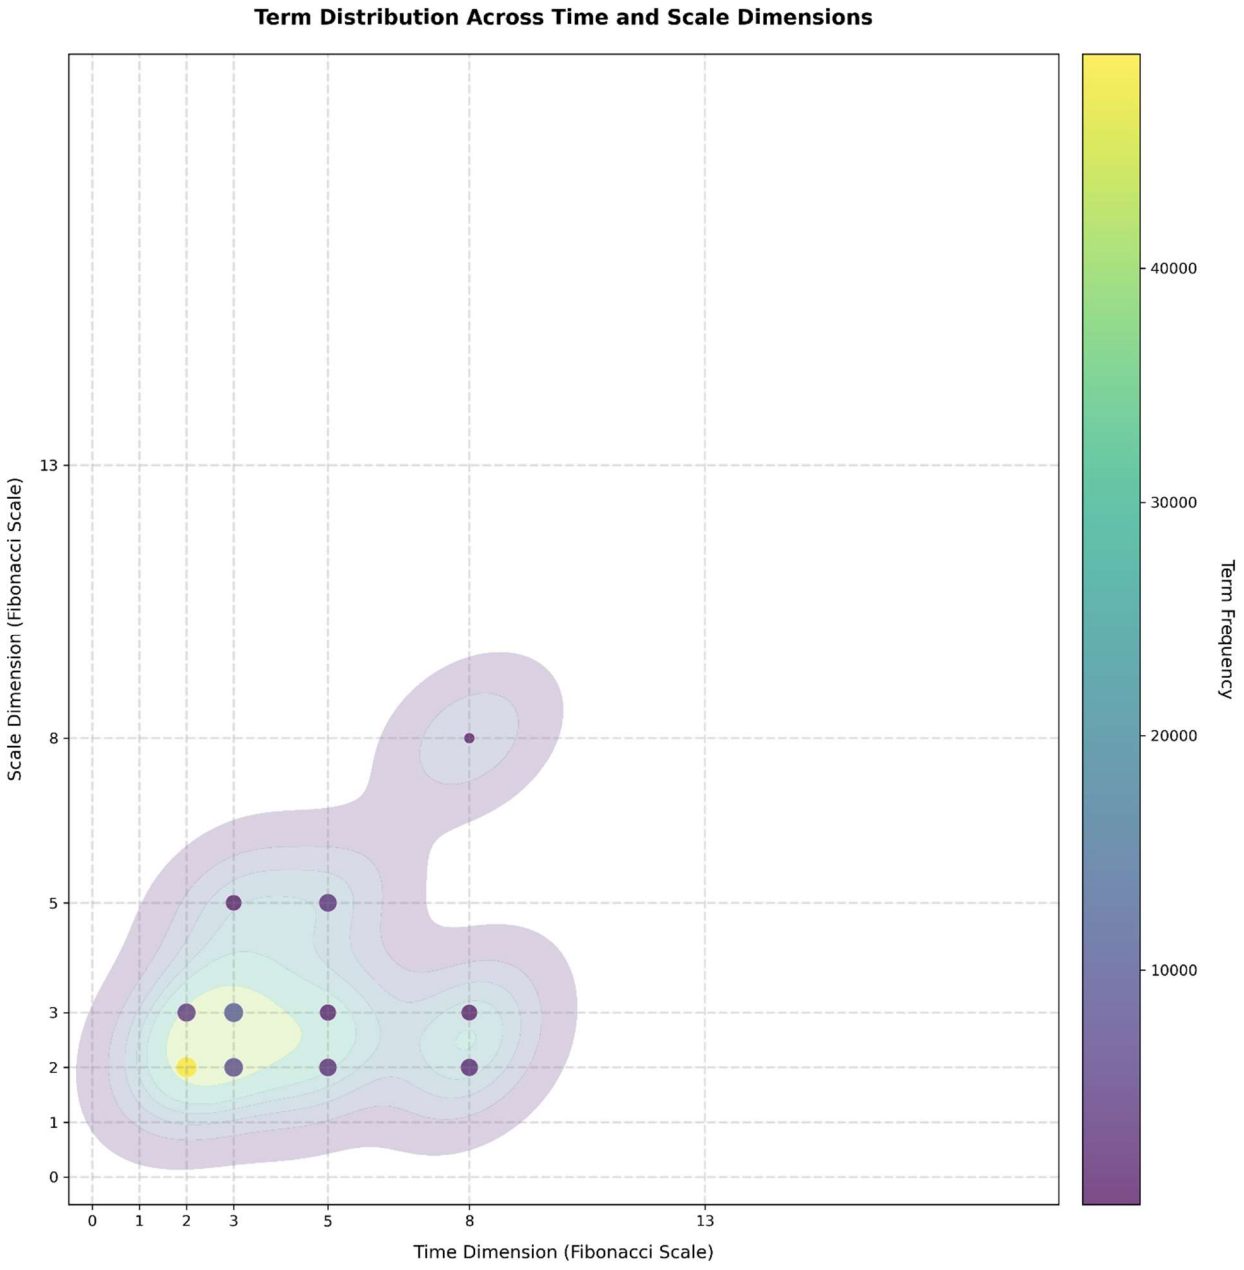

Figure 32: Contour plot showing the classification of resilience features across time and scale for 37 cyber texts.

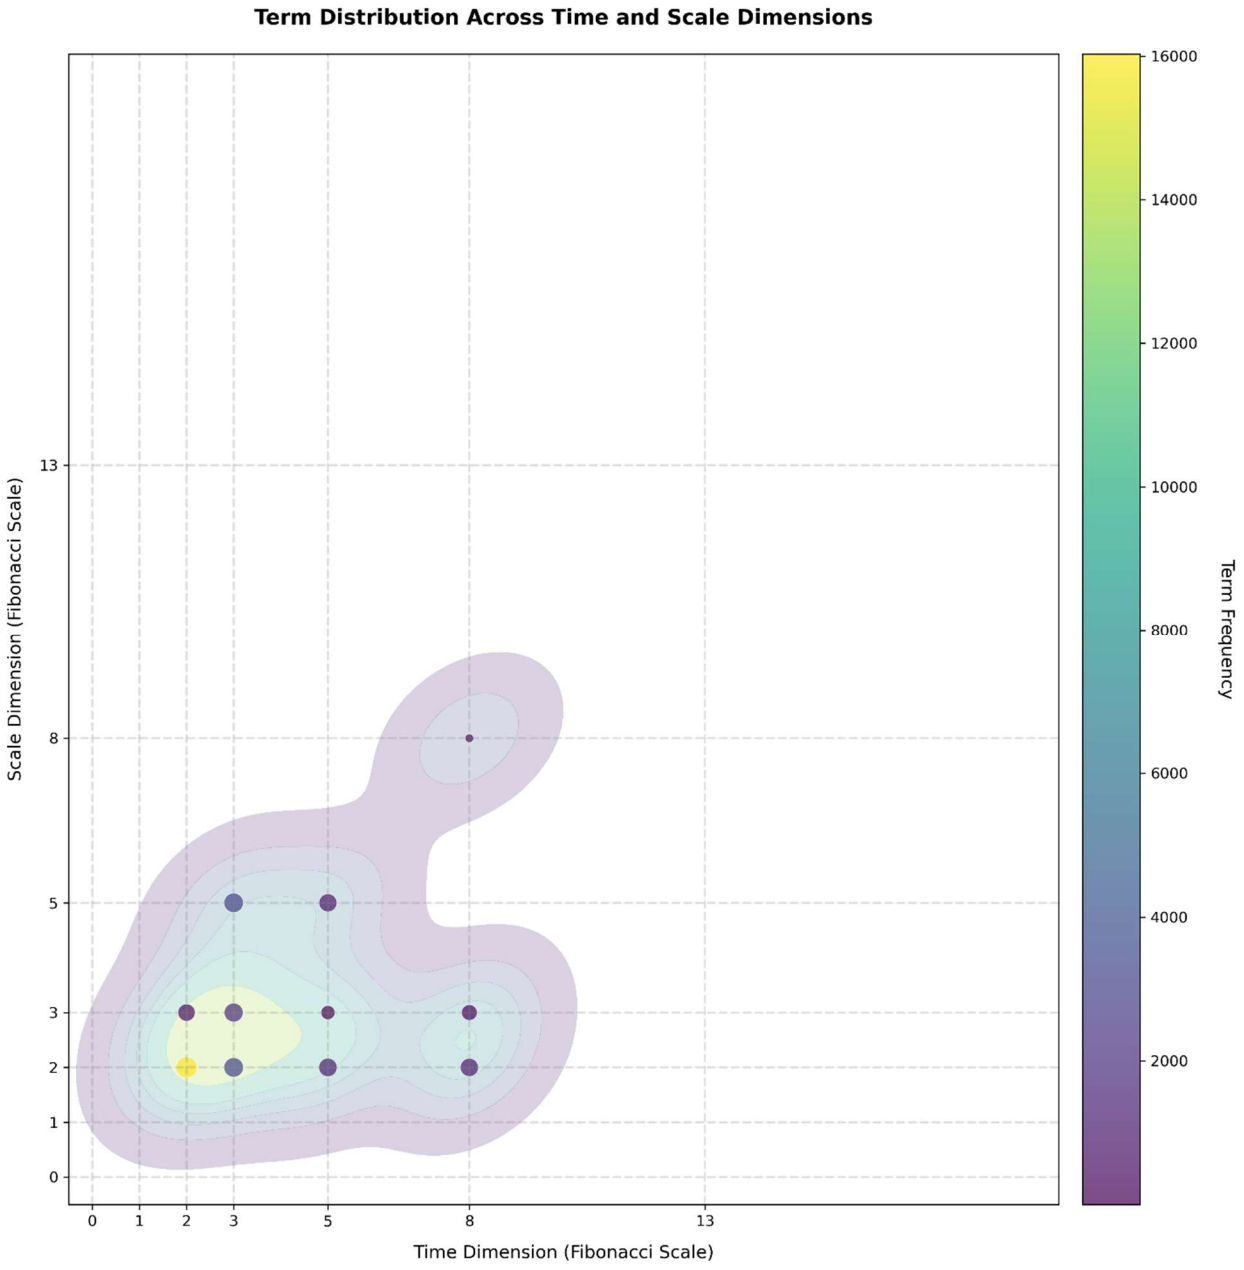

Figure 33: Contour plot showing the classification of resilience features across time and scale for twelve non-cyber texts.

Numerical Results from the Classification Algorithm

From the classification, the mean time and scale values for each document can be calculated, along with the mean values for the cyber, non-cyber, and full corpus of documents. Given the plot densities shown in Figure 31, the mean values are somewhere between 2-3 for time and scale. Figure 34, below, shows the data for these mean values, with the accompanying standard deviations, in box plot format. The mean values for the three sets were between 2.65-2.73 for the time parameter and 2.33-2.68 for the scale parameters.

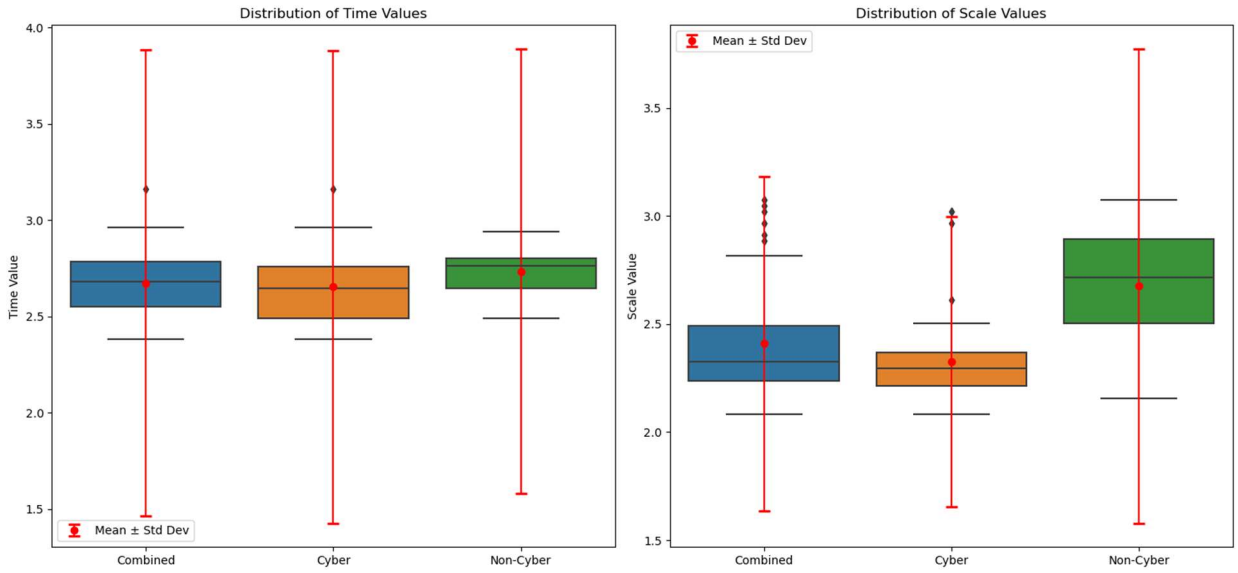

Figure 34: Box plots of the mean time and scale values from the classification algorithm, with their standard deviations, for the combined, cyber, and non-cyber documents.

The non-cyber texts had slightly higher mean values, but the results are well within the standard deviation for the combined corpus for the time values to be not statistically significant. There are significant changes in the scale values and these can be attributed to the fact that those twelve documents were primarily targeting audiences at the community/sector to national-international levels, so that the language would necessarily be classified into that range over the technical and sociotechnical range.

For additional clarity on how the two sets of documents classified, Figure 35, below, shows the plotted mean time and scale values for each document, coded by whether the document was cyber or non-cyber.

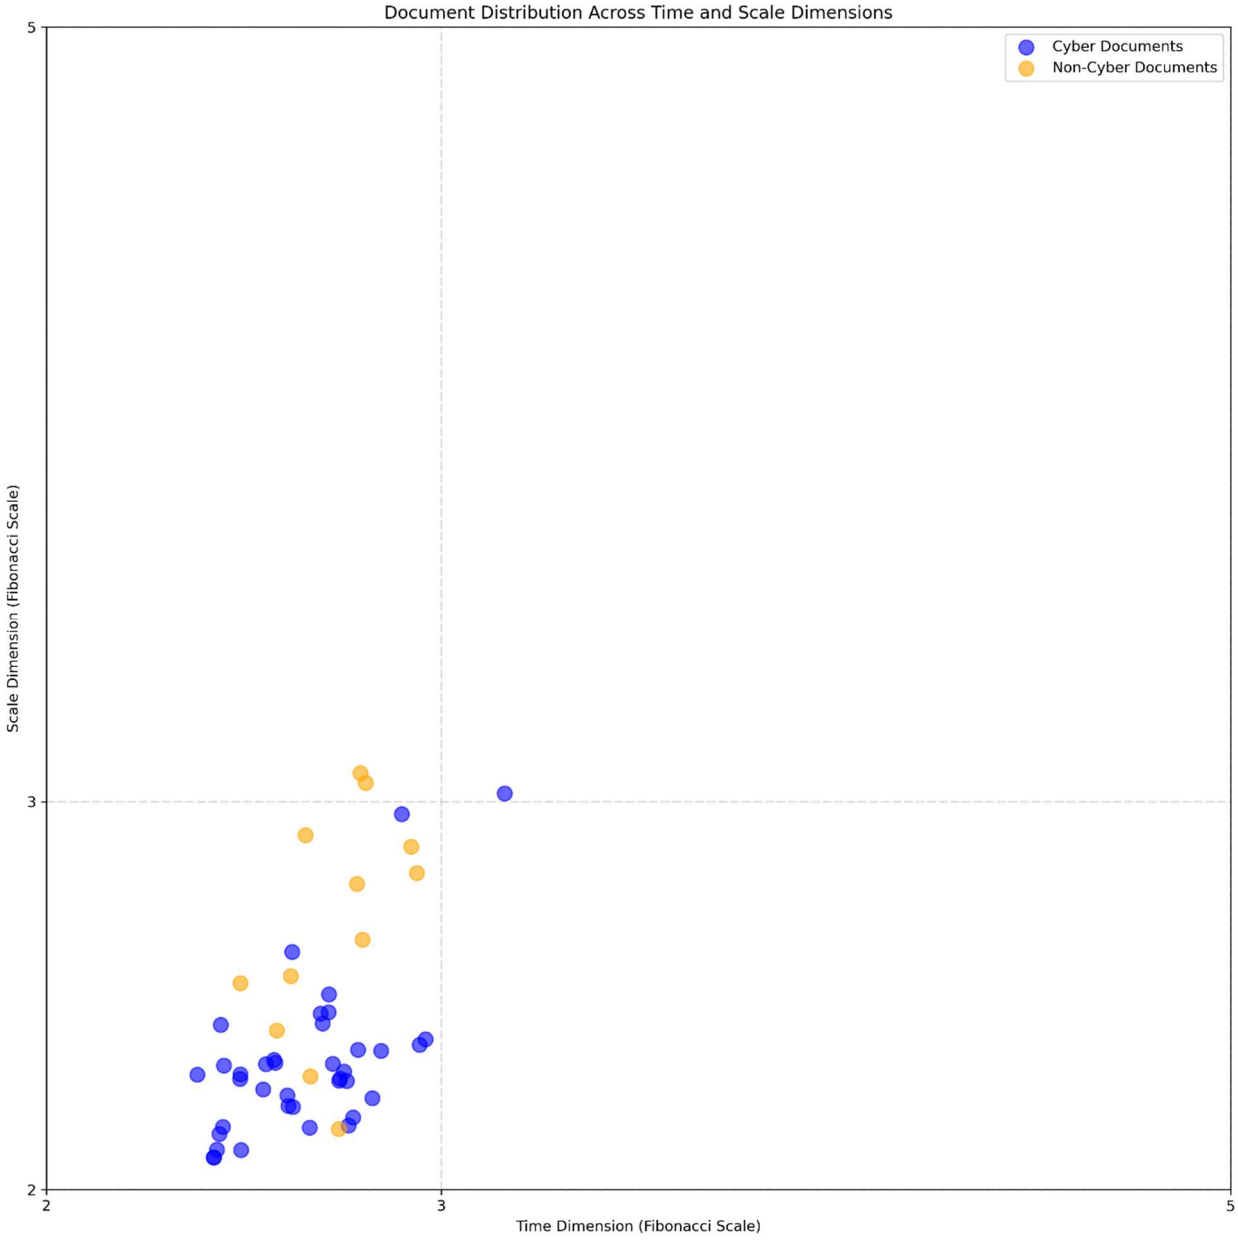

Figure 35: Scatter plot of the mean time and scale values for each of the 49 texts, coded by color for inclusion in the cyber or non-cyber texts groupings.

The zoomed in graphic shows the tight grouping of the mean time and scale values for the distribution of documents. As is visually apparent, there is no significant clustering between the cyber and non-cyber texts, as the markers are commingled.

Documents with the Highest Mean Time and Scale Values

Intuitively, the higher the mean value for time or scale for a given document, the broader the coverage across all levels of time or scale. The combined plot of all documents, Figure 31, reveals that the corpus does not have any coverage at the highest levels of time or scale, which is understandable and expected. The contour indicates that the majority of the features in the combined corpus are primarily at the lower ends of the scales. For the cyber and non-cyber subgroupings of documents, the two documents that had the highest mean time and scale values, and thus the most likely to contain guidance that might satisfy Hypothesis 1, were:

- Draft NIST Special Publication 1800-35 Implementing a Zero Trust Architecture [149]
- Community System Resilience Initiative Steering Committee Final Report [176]

An in-depth, qualitative analysis of each document would be necessary to identify which topics and features the algorithm extracted that drove the classification to the maximum mean values from the corpus. But from the contour plots of both documents, in Figure 36 and Figure 38, it is clear that while the preponderance of features are at the technical to sociotechnical and short time scales, there is a greater proportion of features that classify at higher values of both time and scale, especially at the organizational level for scale and mid-term time level. A word cloud analysis for both documents, in Figure 37 and Figure 39, shows the top fifty terms in each text,

providing a rapid look at whether the language is different enough to drive the higher mean scores.

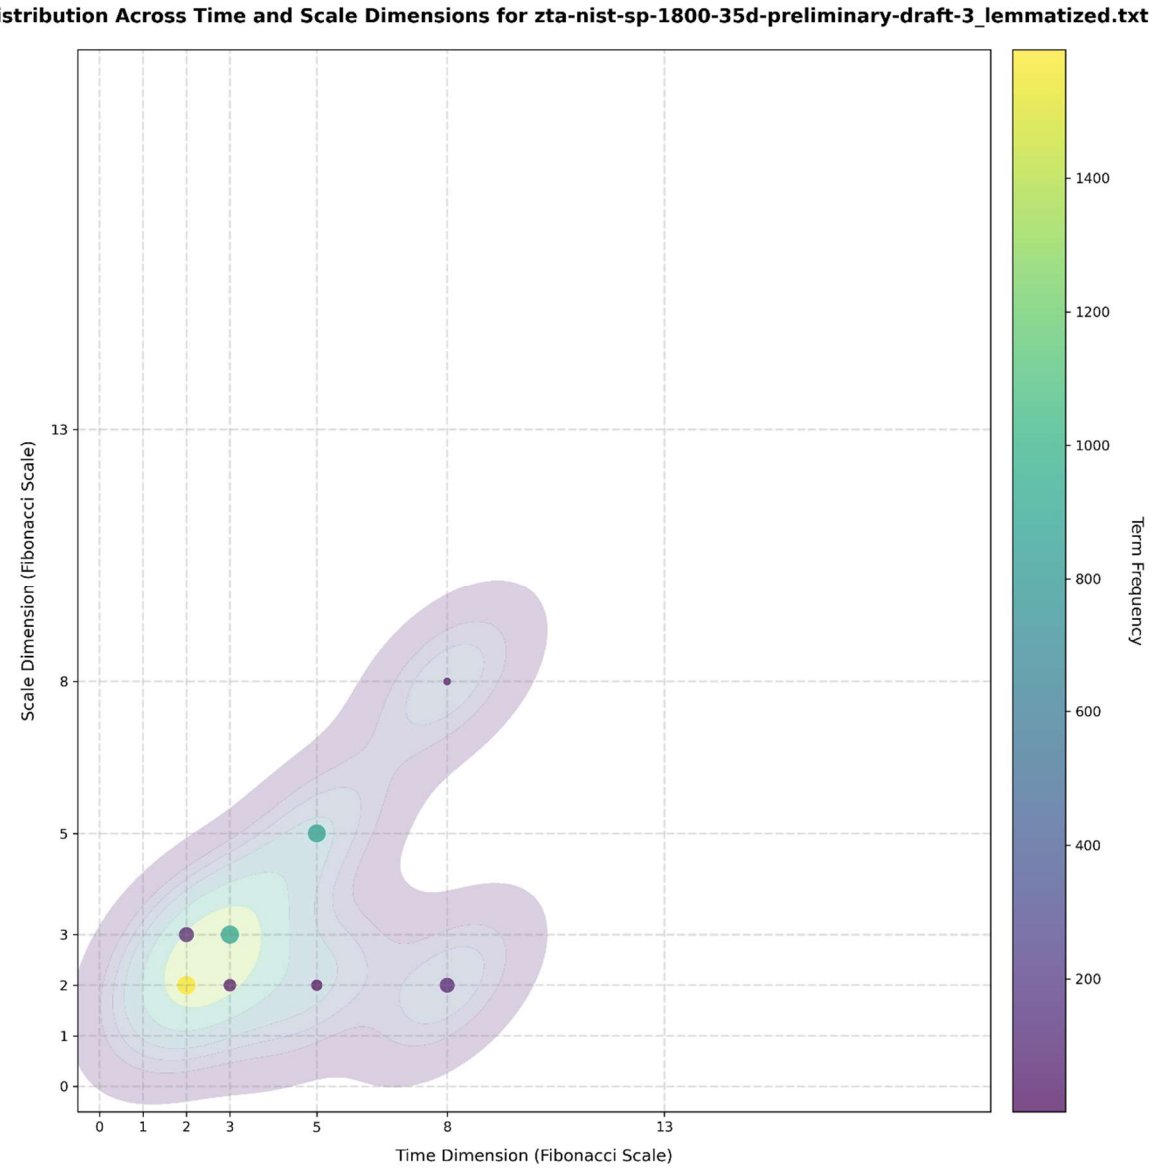

Figure 36: Contour plot of the highest scoring cyber document from the classification algorithm, the draft NIST Special Publication 1800-35 Implementing a Zero Trust Architecture [149].

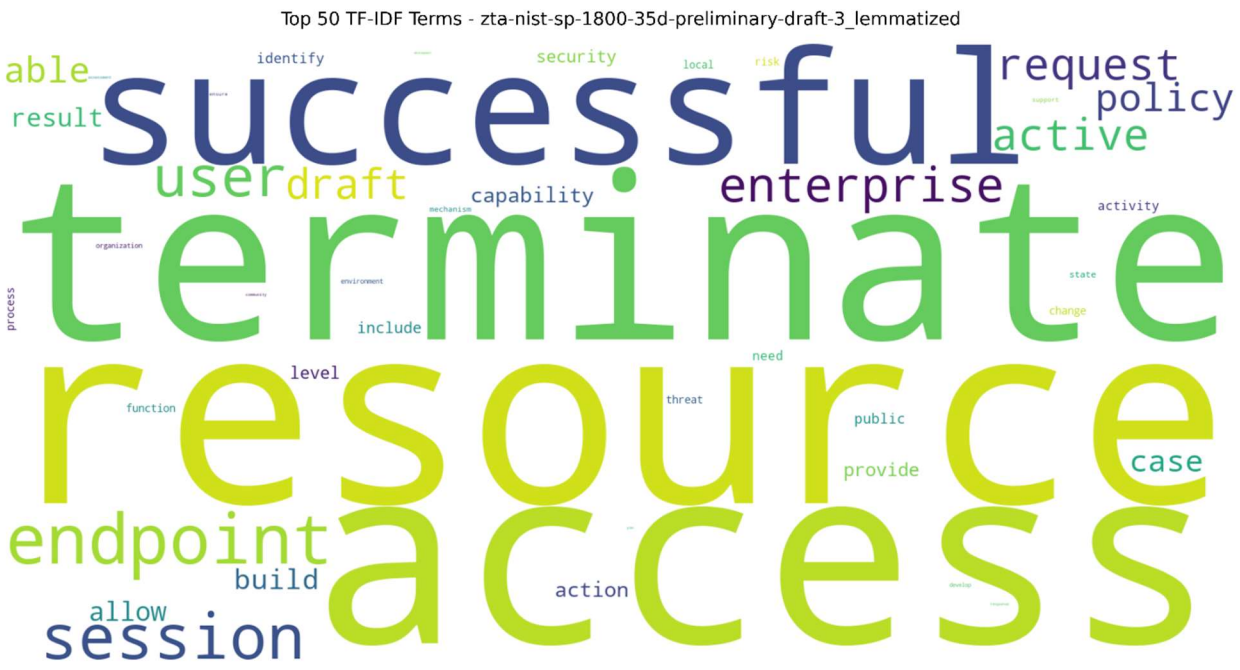

*Figure 37: Word cloud analysis of the top fifty terms from a tf*idf analysis on the draft NIST SP 1800-35 Implementing a Zero Trust Architecture [149].*

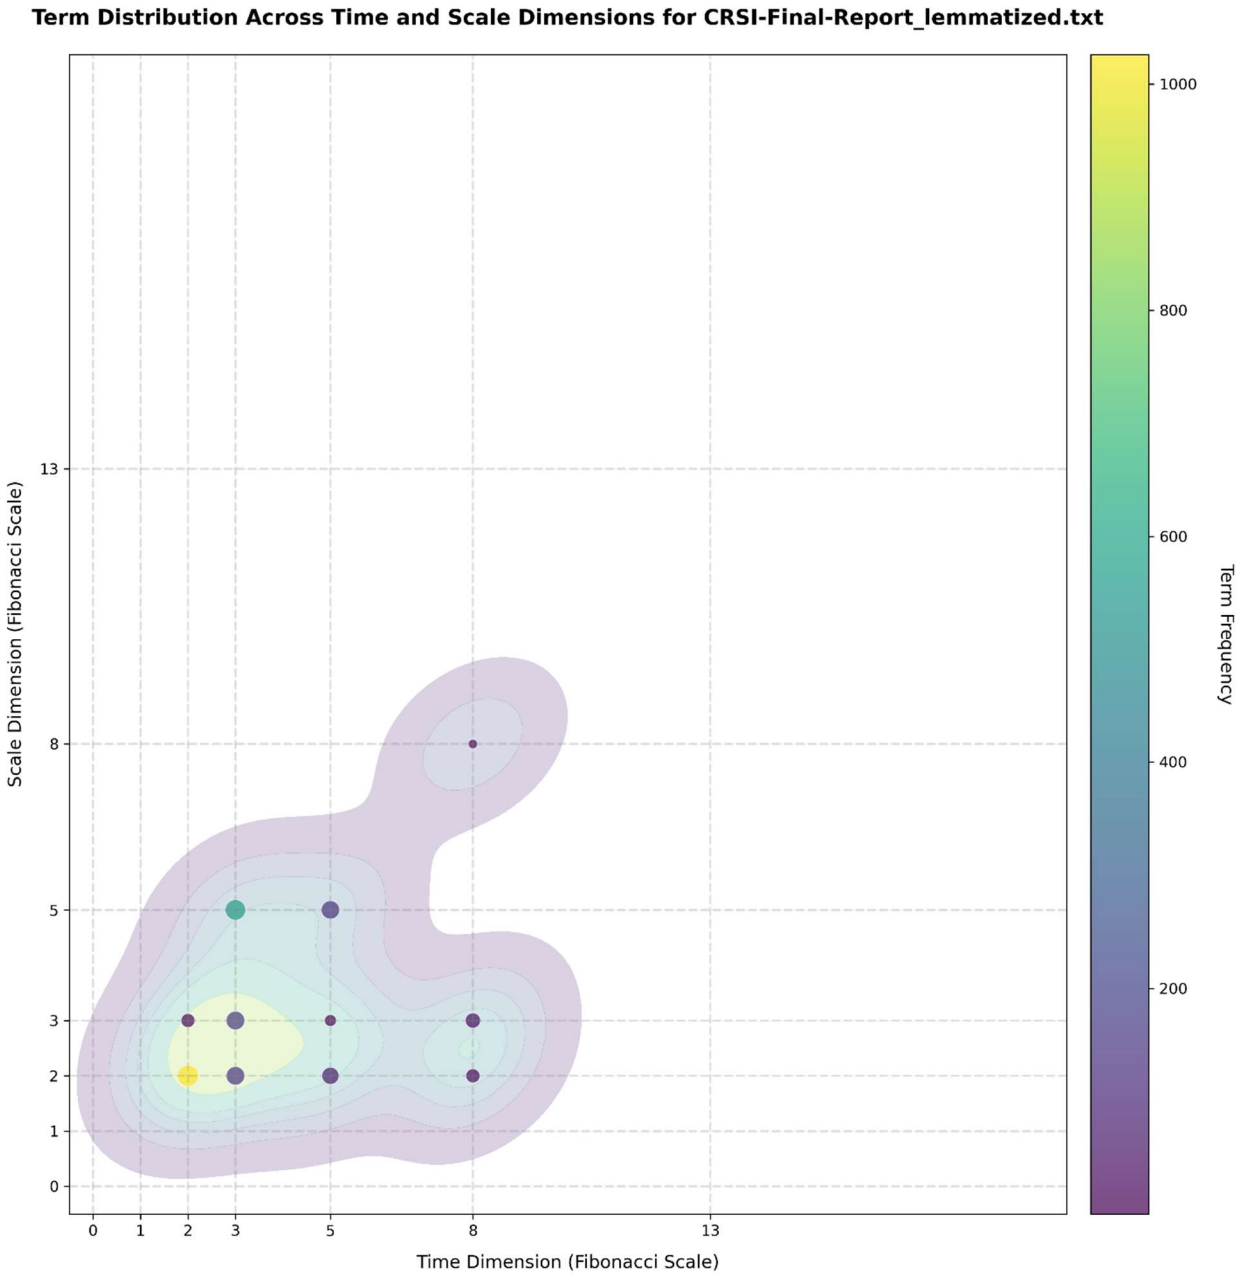

Figure 38: Contour plot of the highest scoring non-cyber document from the classification algorithm, the Community System Resilience Initiative Steering Committee Final Report [176].

*Figure 39: Word cloud analysis of the top fifty terms from a tf*idf analysis on the Community System Resilience Initiative Steering Committee Final Report [176].*

As Figure 39 indicates, the language is closer to the terms identified in Chapter 3, and markedly different from the language used in the cyber texts analyzed in Chapter 4. Figure 37, however, reveals that the high mean score is likely, however, reveals that the high mean score is most likely a flaw in the classification algorithm—a false positive, per se—rather than the text being fundamentally different from the other cyber texts in terms of its cyber resilience content. The top fifty terms in that text closely mirror the technical language in the other cyber texts, as identified in Chapter 4. Zero trust principles are a key concept for developing robust systems and resilience, but NIST SP 1800-35 does not appear to offer guidance at all levels of time and scale for improving cyber resilience based on this analysis.

Comparison to Explicit Cyber Resilience and National Resilience Strategies

This final section considers several documents that explicitly identify cyber resilience or national resilience as their stated objectives. While the combined contour plot in Figure 31 includes the results of those documents, the analysis here is meant to reinforce the conclusions thus far with respect to the primary research question and Hypothesis 1: does a single document exist to provide organizations at multiple levels of society with support for improving cyber resilience? These three documents appear to be the closest, colloquially, to meeting that intent:

- World Economic Forum's Cyber Resilience Index [164]
- National Resilience Strategy of the United States [170]
- Ukraine's strategy on National Resilience in a Changing Security Environment [171]

The set of contour plots and word clouds for these documents follow.

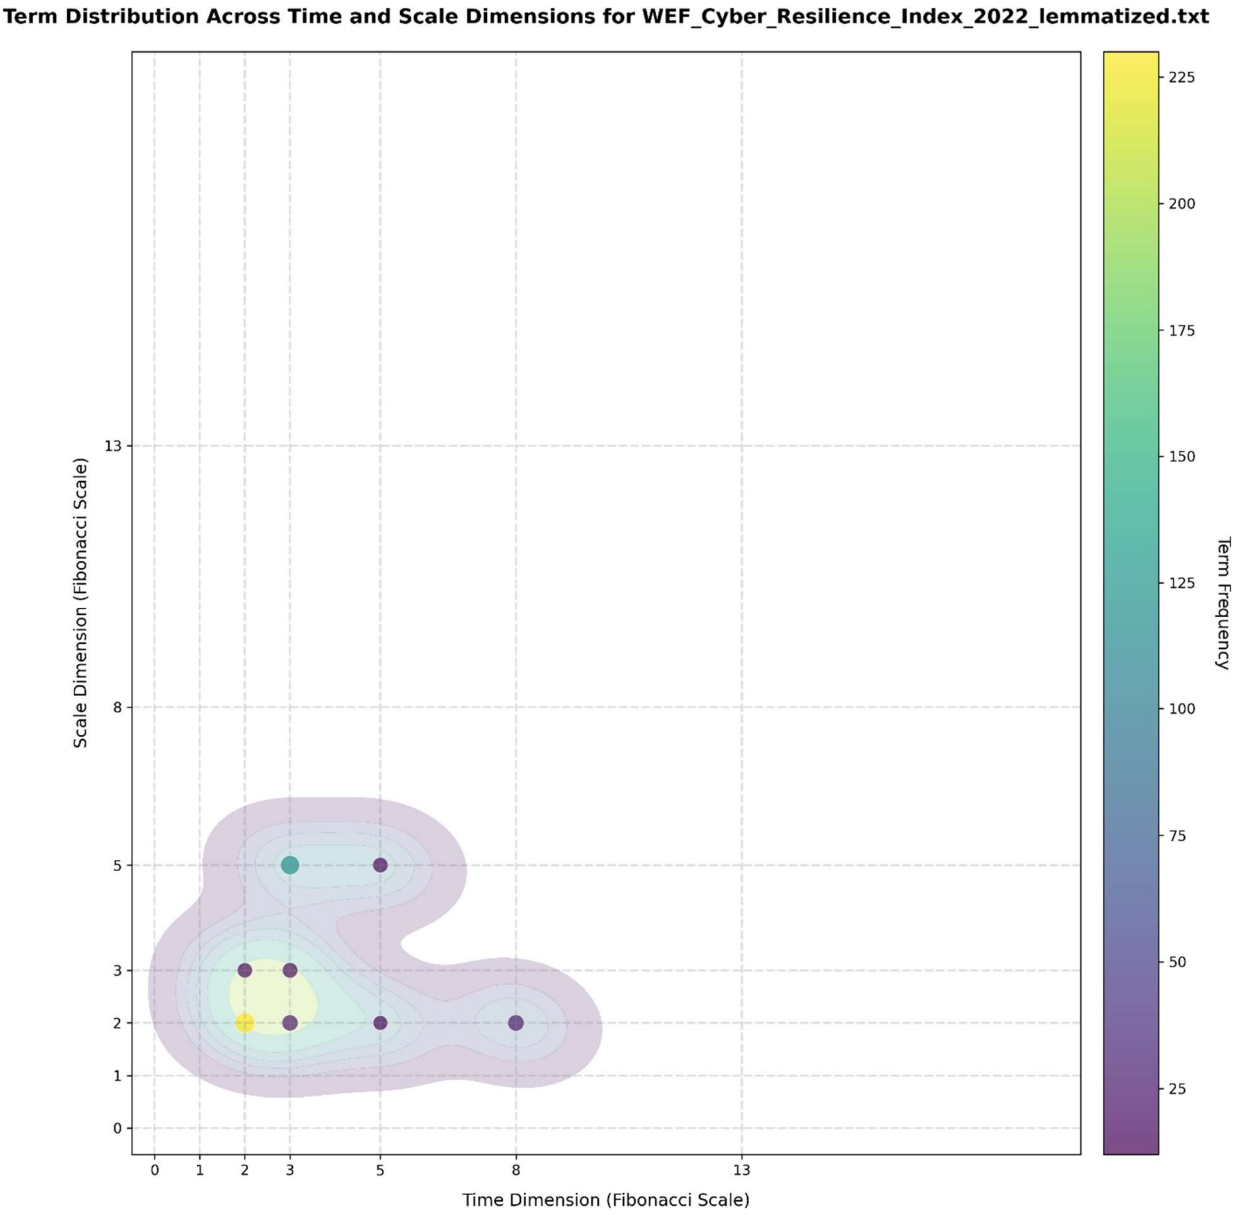

Figure 40: Contour plot from the classification algorithm for the World Economic Forum's Cyber Resilience Index [164].

*Figure 41: Word cloud analysis of the top fifty terms from a tf*idf analysis on the World Economic Forum's Cyber Resilience Index [164].*

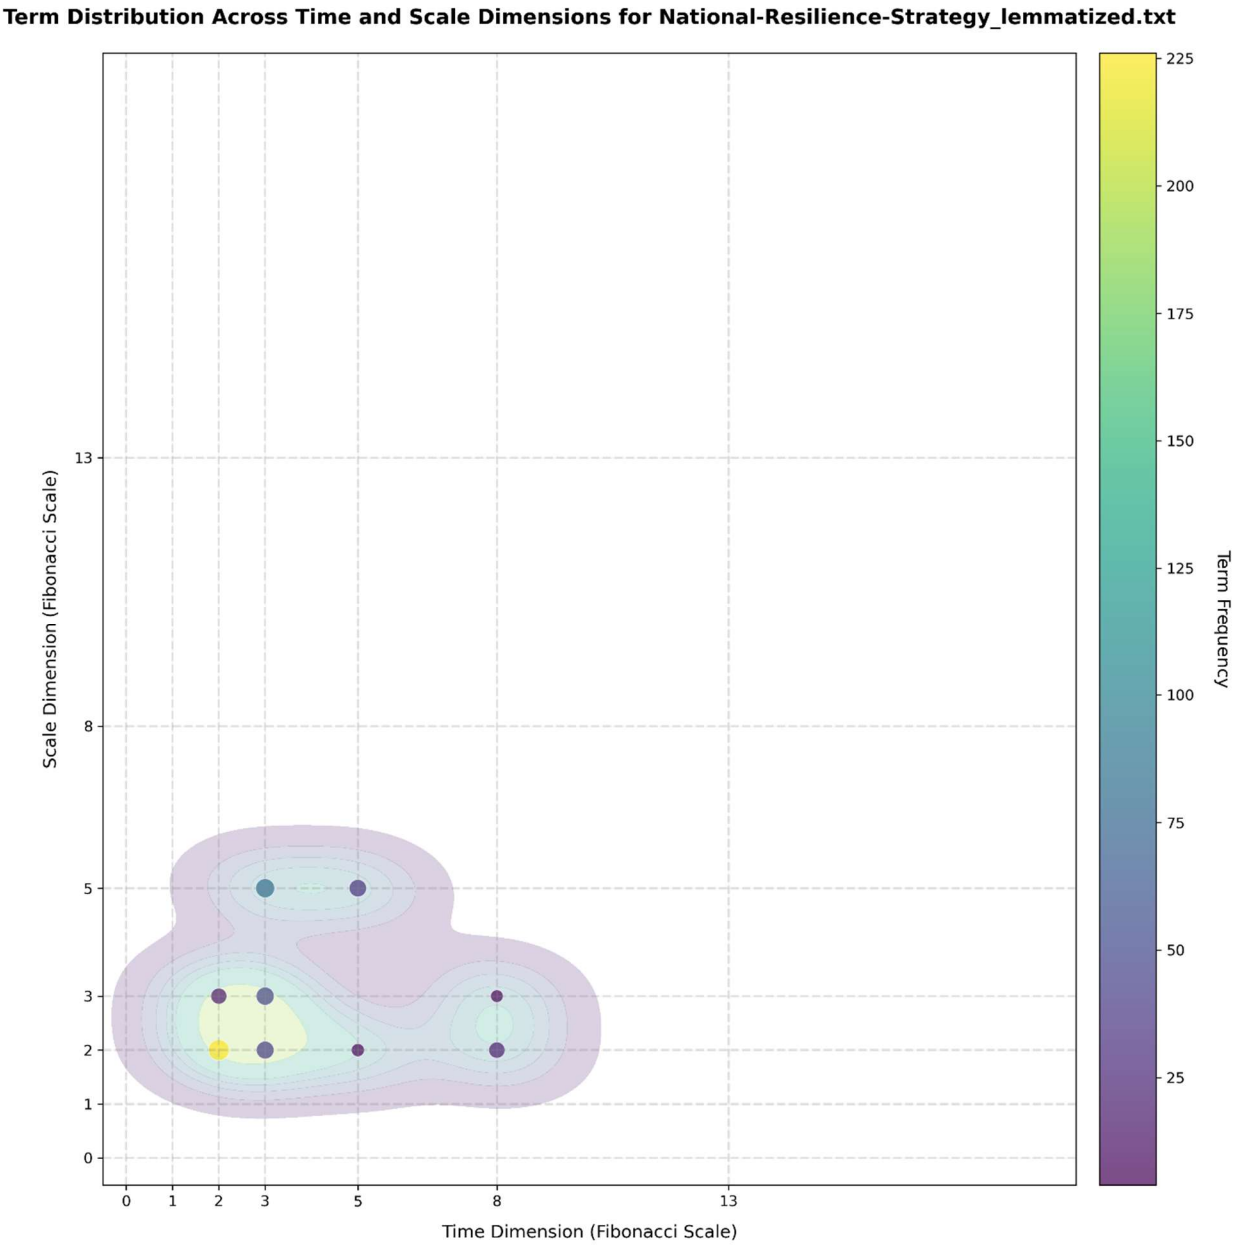

Figure 42: Contour plot from the classification algorithm for the National Resilience Strategy of the United States [170].

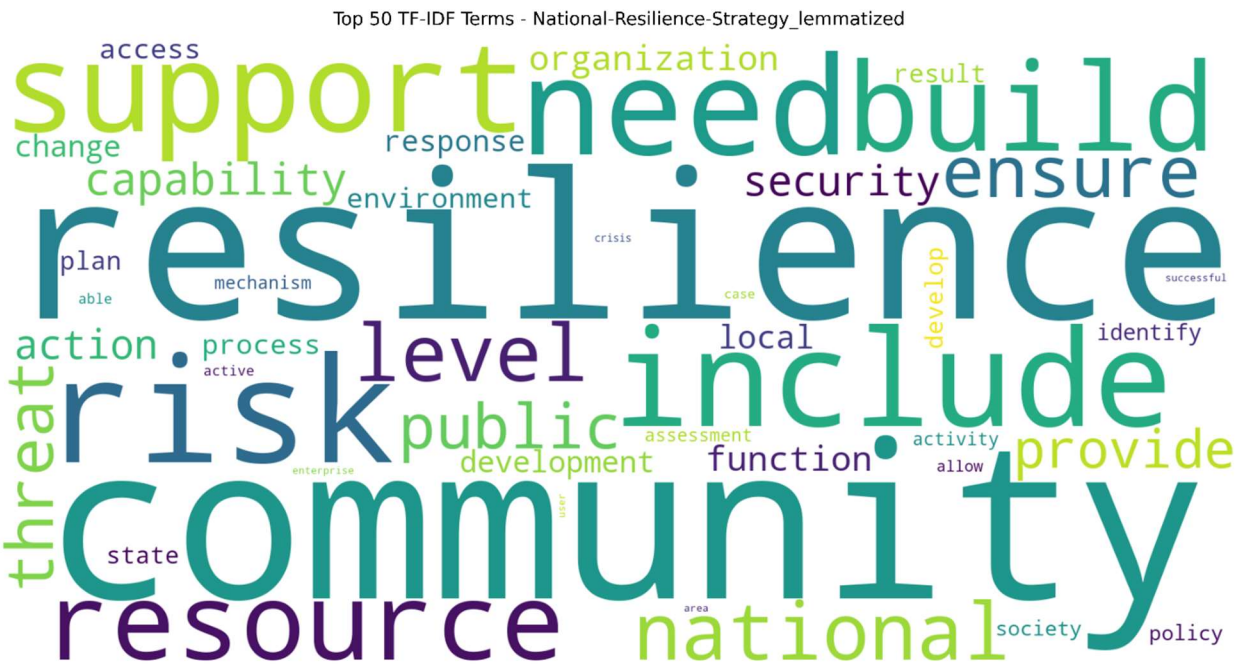

*Figure 43: Word cloud analysis of the top fifty terms from a tf*idf analysis on the National Resilience Strategy of the United States [170].*

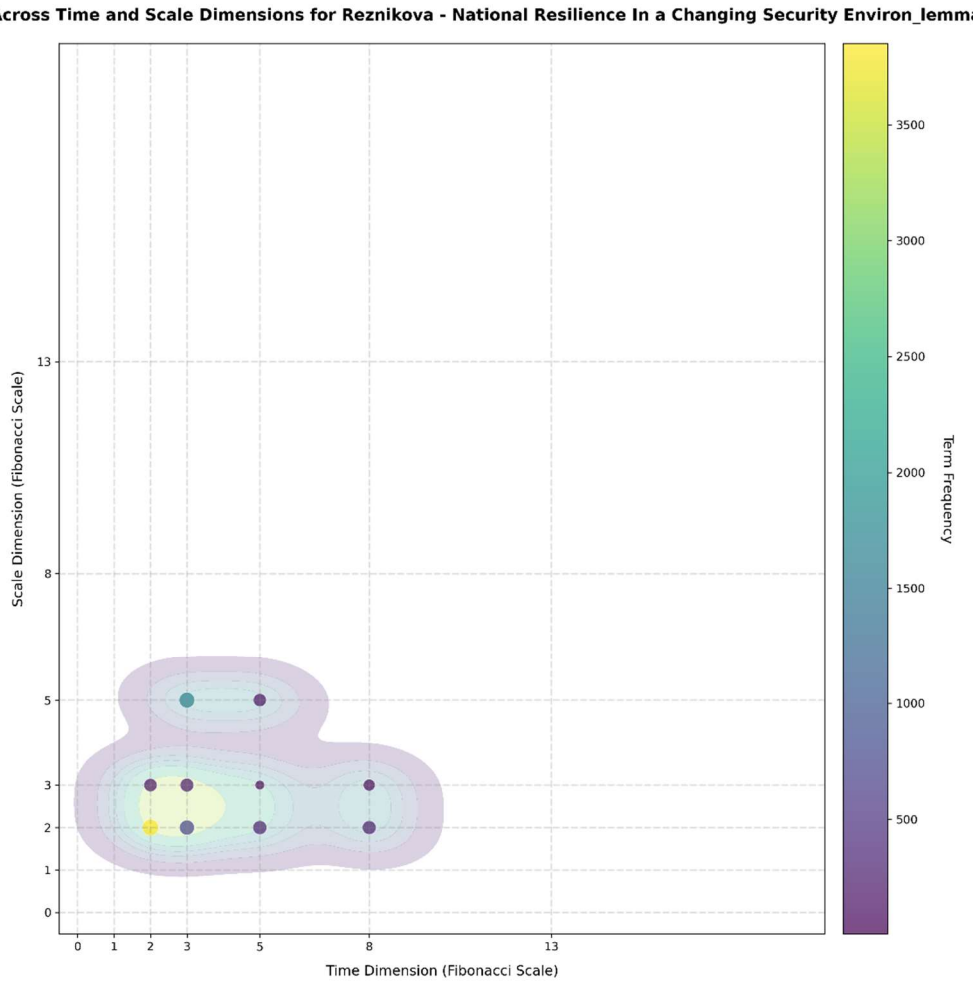

Figure 44: Contour plot from the classification algorithm for Ukraine's National Resilience in a Changing Security Environment [171].

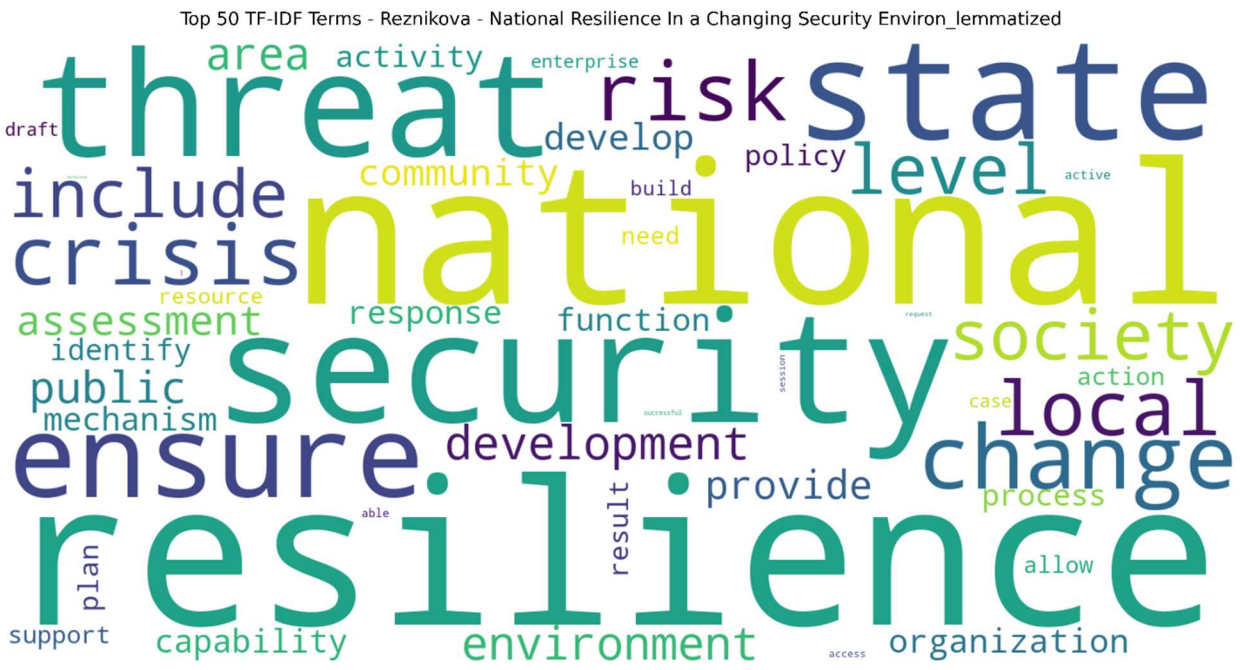

*Figure 45: Word cloud analysis of the top fifty terms from a tf*idf analysis on Ukraine's strategy on National Resilience in a Changing Security Environment [171].*

All three texts show a better distribution of features across multiple scales compared to the body of cyber texts, with greater frequencies of topics at the organizational or community levels for scale and the longer time scales, emphasizing a broader reach of the strategies and greater guidance. However, the national strategies of the United States and Ukraine are not meant to provide individual organizations or communities with guidance to improve resilience; instead, they are meant to align national priorities and resources to achieve national-level resilience outcomes. The World Economic Forum's Cyber Resilience Index qualitatively comes the closest to meeting the objective sought by the research question, but a more thorough analysis of that document would be needed to assess it for improvements to address all aspects of time and scale for improving cyber resilience.

CHAPTER 6: CONCLUSIONS AND FUTURE RESEARCH

Conclusions

With software serving as a central enabling technology that governs how we work, deliver value, and manage our lives, cyber defenders are tasked with designing, integrating, provisioning, and protecting systems that facilitate these outcomes. The growing number of users, both internal and external, that interact with those systems has expanded exponentially in recent decades. Reports of data breaches and other cyberattacks erode our trust in these systems. The guidance available to the cybersecurity community offers best practices for securing systems, planning for and responding to incidents, managing user trust and access, and numerous other topics. However, this guidance does not extend into the sociotechnical domain to help organizations understand how they should be operating to improve their cyber resilience.

In the global cyber threat environment, organizations in all sectors must operate without setbacks. However, the cybersecurity industry's language can blur the lines among terms such as security, resilience, reliability, and robustness. This research conducted an extensive, interdisciplinary study of resilience to determine the central concepts that define it using statistical modeling tools and machine learning algorithms. The topical results from the Latent Dirichlet Allocation analysis (Table 4) provide a concise summary of what the interdisciplinary community defines as resilient:

- Resilience is related to the system’s capability to handle stress and experiences in a positive light

- Resilience is related to the process a system uses to respond to stressful events over a time scale
- Resilience is related to disruptive, impactful, and other extreme events and the ability to survive those events
- Resilience is related to the ability and capacity for a system to change in response to an event
- Resilience is related to a system's functions across scales or a community (vice individual or organizational) in response to crises or stressors
- Resilience is related to a system's capability or ability to respond to a disturbance or stressor
- Resilience is related to the ability of a system to recover within a given state
- Resilience is related to a system's capacity to absorb a disturbance, continue functioning, and change in response
- Resilience is related to mitigating the effects of stressors through social capacities and the environment
- Resilience is related to a system's ability to respond positively under variable conditions
- Resilience is related to a system's ability to adapt in response to specified or particular shocks

Thus, resilience cannot be easily defined with a single definition. Instead, it is best understood as a set of concepts that describe overarching emergent behavior. That behavior is highly dependent on the overall outcomes that the organization seeks—in a sense, it is a type of optimization problem. Biasing to be resilient toward one set of outcomes may make the organization brittle to

other crises that it is not primarily oriented to receive. Thus, general resilience may be a fleeting goal.

Summary Review of Research Questions and Hypotheses

Couched by the interdisciplinary study into resilience, this research sought to answer two primary questions: First, does a single strategy or guidance document exist to provide organizations at multiple levels of society with support for improving cyber resilience? Second, how do those existing documents address resilience? This generated four hypotheses with associated falsification criteria, and the analyses produced the following results.

- Hypothesis 1 (H1): If no single document exists that addresses most or all aspects of resilience, then organizations will not have sufficient guidance to develop complete strategies to improve cyber resilience.
    - Falsification 1 (F1): A document exists that addresses most or all aspects of resilience as identified in the work to test Hypothesis 3 (H3).
    - Result: From the results in Chapters 4 and 5, summarized by Figure 31, no document exists that addresses most or all of the aspects of resilience. Thus, we can reject the Falsification 1 criteria and accept Hypothesis 1 as true.
- Hypothesis 2 (H2): If one or more documents exist that address some aspects of resilience, then they can be leveraged and expanded to create a single document for organizations to improve cyber resilience.
    - Falsification 2 (F2): Hypothesis 1 (H1) is proven true.

- Result: Since we accept Hypothesis 1 is true, that satisfies the Falsification 2 criteria. We must reject Hypothesis 2. No combination of documents exist that can be leveraged to create a single document for organizations to improve cyber resilience.

- Hypothesis 3 (H3): If the existing documents incorporate most or all aspects of resilience, then there should be significant overlap with the interdisciplinary research and concepts on resilience.
  - Falsification 3 (F3): The analysis indicates gaps between the interdisciplinary concepts on resilience and the existing documents.
  - Result: The results of the tf*idf analysis, Latent Dirichlet Allocation analysis, and similarity analyses in Chapters 3 and 5 comparing the documents with the interdisciplinary concept of resilience validates the Falsification 3 criteria—gaps do exist between the two sets of work, which is ultimately evidenced in the acceptance of Hypothesis 1.
- Hypothesis 4 (H4): If the existing documents focus primarily on the technical and sociotechnical aspects of cybersecurity, then the documents will lack sufficient guidance on most or all aspects of resilience.
  - Falsification 4 (F4): The analysis indicates few to no gaps in one or more existing documents.
  - Result: The mean time and scale values for the 37 cyber texts and twelve non-cyber tests (the box plots in Figure 34) and the various contour plots presented in Chapter 5 all indicate a strong bias toward the technical and sociotechnical ends of the time and scale axes. We can accept Hypothesis 4.

In summary, the acceptance of Hypotheses 1 and 4 and rejection of Hypotheses 2 and 3 in this research leads to the conclusion that the existing body of cybersecurity documents, even ones that seek to provide guidance on fostering cyber resilience to organizations, lack a sufficient connection to the interdisciplinary conceptualization of resilience. Further developments will be needed to better address the aspects of resilience across time and scale to provide organizations with the necessary information and tools to develop strategies to improve their cyber resilience.

Contributions

This research makes several contributions in the areas of research on cyber resilience, practical contributions that organizations can draw on, and theoretical contributions from the methodology used. As stated in Chapter 1, this research expected to make the following three primary contributions toward cyber resilience and did so with the specified results.

1. A classification framework from which future frameworks or guiding documents can be assessed for features of resilience—cyber or otherwise.
    a. Result: The development of a classification algorithm based on statistical modeling and machine learning techniques provides an algorithm that can be used to assess future cyber security or other resilience texts. It also provides a basis for continued development and refinement of the algorithm through expanding the dictionary and better or finer grained measurements for time and scale properties. The utilization of a standard data curation pipeline approach simplifies implementation on future texts.

2. An assessment of the existing frameworks and the extent to which the framework addresses some or all aspects of resilience as the interdisciplinary community understands it.
    a. Result: This research assessed 37 cyber security related texts encompassing narrow technical publications to sector- or industry-specific frameworks promoting resilience and found that no single framework adequately addresses resilience as the interdisciplinary community understands it. Of note, this research also analyzed twelve non-cyber, but resilience-related texts and found gaps in coverage from the interdisciplinary community's understanding of resilience.
3. Identify the gaps in the frameworks, or combined set of frameworks, to show where the cybersecurity community can focus future efforts on developing guidance to support all aspects of resilience, thus giving organizations the ability to develop strategies to improve cyber resilience.
    a. Result: Chapter 5 showed that the cybersecurity texts, as expected, primarily focus on the technical and sociotechnical aspects of cybersecurity, with insufficient to no coverage on some aspects of resilience.

Practical Contributions for Organizations

From a practical standpoint, organizations can use the results of this research to reframe resilience objectives and improve cyber resilience strategies. While not explicitly stated or oriented as such, this research makes three contributions that practitioners can leverage to improve organizational cyber resilience.

1. A coding scheme that can be used with a body of organizational documentation to assess the current language against the resilience and cybersecurity dictionaries, potentially providing a quantitative measure that can be used as a basis for measuring improvement.
2. Identification of areas where existing cybersecurity guidance falls short of meeting resilience objectives, giving organizations the knowledge that their efforts may not be fully meeting expectations.
3. A detailed development of what resilience is, and the set of behaviors or attributes that define resilient systems. The prior literature focuses on a singular definition of resilience, which can be difficult for practitioners or organizations to easily adapt to their specific situation or ecosystem. The decomposition of resilience from a single definition approach into a collection of attributes gives practitioners greater granularity, from which they can better understand how their organization measures up. There is future work in this area to further explore the connections with related concepts, such as high reliability organizations, system safety, organizational culture, etc. and how those bodies of work might link to and enhance organizational resilience.

For the broader cybersecurity community, the natural progression of this research would be to begin mapping the controls in the 37 cybersecurity frameworks into a broader scaffolding around the major themes present. By aligning the controls with the attributes of resilience more narrowly, the community will be able to determine where existing controls meet the desired outcomes. The technical and sociotechnical controls, as presented in this research, ensure that those domains would be well-represented in this broader scaffolding, whereas the organizational controls, governance, risk, and compliance, and integration with the broader business strategy

and value proposition areas would likely need attention. In recent years, corporate boards have begun to realize their role in cybersecurity and its connection to delivering business value, but the debate on the proper role of the board, its ability to link business outcomes to cybersecurity lines of effort, and define the desired resilient attributes, such as the critical functions of the business and acceptance of certain thresholds for providing value, are still being debated. Fundamental to this debate is understanding what resilience is, determining how the organization would like to define its desired resilient outcomes, and understanding how to link existing cybersecurity processes to resilient outcomes.

This research sought to provide the foundations of that next step, since the classification of existing cybersecurity controls onto the classification framework and identification of new resilience controls or practices must be explicitly tied to the organization's desired resilience outcomes. Broad research across the cybersecurity and interdisciplinary communities is needed to develop a core set of resilience controls that most organizations can adopt, akin to what each of the cybersecurity texts, such as the NIST Cybersecurity Framework 2.0, sought to provide for security purposes. Development of these controls, as has been the case for development of cybersecurity standards, will involve much learning, debate, and experimentation to determine which controls are more effective. This will likely require the explicit bridging or integration of cybersecurity with traditional business domains, viewed through the lens of resilience.

Theoretical Contributions

Finally, this research produced several theoretical contributions that are topical or methodological. These serve the research community by providing new insights or methods. This research made four primary theoretical or methodological contributions.

1. The use of statistical modeling and machine learning algorithms on the interdisciplinary body of knowledge on resilience appears to be novel. The extensive literature review conducted to develop the resilience dictionary did not uncover any existing research that used this approach. The topical results from the Latent Dirichlet Allocation analysis, for example, summarize a large body of interdisciplinary research in a way that is not present in the existing literature. This approach may be useful in other areas of research where there is a substantial interdisciplinary component.
2. This research expanded and integrated the concept of resilience by demonstrating that efforts toward a universal definition may not be as beneficial as coalescing around the set of attributes associated with resilient systems. The literature review did uncover several meta-studies on resilience, but the majority of these studies were focused on a singular definition with qualitative analyses.
3. The development of a novel classification framework using the same statistical modeling and machine learning approaches with a data curation pipeline to conduct large scale semantic analysis of texts appears to be novel. While natural language processing and other semantic analyses are in widespread use, the literature review and associated research for this dissertation did not uncover any similar approaches to analyzing a corpus.

4. This research contributed to a broader understanding of the relationship between existing cybersecurity guidance and literature and the broader corpus of work on resilience, complex adaptive systems, and related concepts. This opens new opportunities for future research.

Future Research

This integrative study opens up multiple opportunities for future research. The integration of an already interdisciplinary field with data science techniques necessarily results in the initial development of new techniques and knowledge and should offer multiple avenues of new research as a result. This final section organizes those opportunities around the theoretical and the practical.

Theoretical Opportunities

1. Investigate and improve the classifiers for the time and scale aspects of resilience against cybersecurity applications. The use of a Fibonacci scale was only a first order approximation to convey non-linearity; there is opportunity to refine the scaling into a more quantitative system with true non-linearity, to include quantification of the adaptive cycles and panarchy.
2. Improve and expand the resilience and cybersecurity dictionary. The initial GPT-aided classification of terms provided a starting point. There is ample opportunity here for qualitative and quantitative research into the classification of those terms.

3. Leverage and connect previous research into quantitative measurements of resilience and cybersecurity to develop this body of research into a quantitative system.
4. Investigation into the "ideal" centroid or mean scores for time and scale values from a classification framework. It is assumed that an unweighted centroid with equal representation at all intersections will not be ideal since there should likely be a bias toward local times and scales, but identification of an ideal centroid position would be of value to the academic and practitioner communities.

Practical Opportunities

1. The large corpus of business-oriented books around human activities, such as strategy, sociotechnical systems, supply chain security and resilience, culture, etc. is ripe for evaluation with the classification framework. This provides the opportunity to link these bodies of knowledge into a coherent whole.
2. Development of a cybersecurity framework that encompasses the majority of the time and scale attributes of resilience. Further, investigation into what types of recommendations or best practices can be truly scale free.
3. Analyzing the roles of the sixteen critical infrastructure sectors identified by CISA with respect to the resilience attributes presented here and how they influence the adaptive cycles above (national to international) and below (sector and below) them [184].

REFERENCES

[1] M. E. Conway, “How do committees invent?,” *Datamation*, pp. 28–31, 1968.

[2] S. Smith, “Towards a Scientific Definition of Cyber Resilience,” *Int. Conf. Cyber Warf. Secur.*, vol. 18, no. 1, Art. no. 1, Feb. 2023, doi: 10.34190/iccws.18.1.960.

[3] S. M. Alhidaifi, M. R. Asghar, and I. S. Ansari, “A Survey on Cyber Resilience: Key Strategies, Research Challenges, and Future Directions,” *ACM Comput Surv*, vol. 56, no. 8, p. 196:1-196:48, Apr. 2024, doi: 10.1145/3649218.

[4] E. Barasa, R. Mbau, and L. Gilson, “What Is Resilience and How Can It Be Nurtured? A Systematic Review of Empirical Literature on Organizational Resilience,” *Int. J. Health Policy Manag.*, vol. 7, no. 6, p. 491, Jun. 2018, doi: 10.15171/IJHPM.2018.06.

[5] J. F. Carías, S. Arrizabalaga, L. Labaka, and J. Hernantes, “Cyber Resilience Progression Model,” *Appl. Sci.*, vol. 10, no. 21, Art. no. 21, Jan. 2020, doi: 10.3390/app10217393.

[6] Y. Y. Haimes, “On the Definition of Resilience in Systems,” *Risk Anal. Int. J.*, vol. 29, no. 4, pp. 498–501, Apr. 2009, doi: 10.1111/j.1539-6924.2009.01216.x.

[7] H. Herrman, D. E. Stewart, N. Diaz-Granados, E. L. Berger, B. Jackson, and T. Yuen, “What is Resilience?,” *Can. J. Psychiatry*, vol. 56, no. 5, pp. 258–265, May 2011, doi: 10.1177/070674371105600504.

[8] A. J. Masys, “The Cyber-Ecosystem Enabling Resilience Through the Comprehensive Approach,” in *Disaster Management: Enabling Resilience*, A. Masys, Ed., in Lecture Notes in Social Networks. , Cham: Springer International Publishing, 2015, pp. 143–154. doi: 10.1007/978-3-319-08819-8_8.

[9] B. Walker, “Resilience: what it is and is not,” *Ecol. Soc.*, vol. 25, Jun. 2020, doi: 10.5751/ES-11647-250211.

[10] J. W. Creswell and J. D. Creswell, *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*. SAGE Publications, 2018.

[11] *ChatGPT*. OpenAI. Accessed: Dec. 10, 2023. [Large Language Model (LLM)]. Available: https://chat.openai.com

[12] *pypdf: A pure-python PDF library capable of splitting, merging, cropping, and transforming PDF files*. Python.

[13] “PyMuPDF 1.24.10 documentation.” Accessed: Sep. 22, 2024. [Online]. Available: https://pymupdf.readthedocs.io/en/latest/index.html

[14] “csv — CSV File Reading and Writing,” Python documentation. Accessed: Sep. 22, 2024. [Online]. Available: https://docs.python.org/3/library/csv.html

[15] W. J. Wilbur and K. Sirotkin, “The automatic identification of stop words,” *J. Inf. Sci.*, vol. 18, no. 1, pp. 45–55, Feb. 1992, doi: 10.1177/016555159201800106.

[16] “NLTK :: Natural Language Toolkit.” Accessed: Sep. 22, 2024. [Online]. Available: https://www.nltk.org/

[17] “re — Regular expression operations — Python 3.12.6 documentation.” Accessed: Sep. 22, 2024. [Online]. Available: https://docs.python.org/3/library/re.html

[18] “spaCy · Industrial-strength Natural Language Processing in Python.” Accessed: Sep. 22, 2024. [Online]. Available: https://spacy.io/

[19] “pickle — Python object serialization,” Python documentation. Accessed: Sep. 22, 2024. [Online]. Available: https://docs.python.org/3/library/pickle.html

[20] "os — Miscellaneous operating system interfaces — Python 3.12.6 documentation." Accessed: Sep. 22, 2024. [Online]. Available: https://docs.python.org/3/library/os.html
[21] "collections — Container datatypes," Python documentation. Accessed: Sep. 22, 2024. [Online]. Available: https://docs.python.org/3/library/collections.html
[22] E. Loper and S. Bird, "NLTK: The Natural Language Toolkit," May 17, 2002, *arXiv*: arXiv:cs/0205028. doi: 10.48550/arXiv.cs/0205028.
[23] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," May 24, 2019, *arXiv*: arXiv:1810.04805. doi: 10.48550/arXiv.1810.04805.
[24] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Inf. Process. Manag.*, vol. 24, no. 5, pp. 513–523, Jan. 1988, doi: 10.1016/0306-4573(88)90021-0.
[25] J. Ramos, "Using TF-IDF to Determine Word Relevance in Document Queries," in *Proceedings of the first instructional conference on machine learning*, 2003, pp. 29–48. [Online]. Available: https://www.researchgate.net/profile/Farshad-Madani/post/In_information_retrieval_tf-idf_calculation_why_we_dont_divide_tf_by_the_length_of_the_related_document/attachment/59d6446679197b807799fae0/AS%3A448525403201536%401483948197307/download/Using+TF-IDF+to+Determine+Word+Relevance+in+Document+Queries.pdf
[26] K. Giles and W. H. Ii, "Divided by a Common Language: Cyber Definitions in Chinese, Russian and English," presented at the 2013 5th International Conference on Cyber Conflict, Tallinn, Estonia: NATO CCD COE Publications, 2013. [Online]. Available: https://www.ccdcoe.org/uploads/2018/10/22_d3r1s1_giles.pdf
[27] *Jacobellis v. Ohio*. 1694. [Online]. Available: https://tile.loc.gov/storage-services/service/ll/usrep/usrep378/usrep378184/usrep378184.pdf
[28] R. Ross, V. Pillitteri, R. Graubart, D. Bodeau, and R. McQuaid, "Developing cyber-resilient systems : a systems security engineering approach," National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST SP 800-160v2r1, Dec. 2021. doi: 10.6028/NIST.SP.800-160v2r1.
[29] "Glossary | CSRC." Accessed: Jul. 11, 2024. [Online]. Available: https://csrc.nist.gov/glossary
[30] "Committee on National Security Systems Glossary." Committee on National Security Systems, Mar. 02, 2022. Accessed: Apr. 14, 2024. [Online]. Available: https://www.cnss.gov/CNSS/issuances/Instructions.cfm
[31] Joint Task Force, "Assessing security and privacy controls in information systems and organizations," National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST SP 800-53Ar5, Jan. 2022. doi: 10.6028/NIST.SP.800-53Ar5.
[32] K. Stouffer *et al.*, "Cybersecurity Framework Version 1.1 Manufacturing Profile," National Institute of Standards and Technology, Oct. 2020. doi: 10.6028/NIST.IR.8183r1.
[33] X. Xue, L. Wang, and R. J. Yang, "Exploring the science of resilience: critical review and bibliometric analysis," *Nat. Hazards*, vol. 90, no. 1, pp. 477–510, Jan. 2018, doi: 10.1007/s11069-017-3040-y.
[34] C. S. Holling, "Resilience and Stability of Ecological Systems," *Annu. Rev. Ecol. Evol. Syst.*, vol. 4, no. Volume 4, 1973, pp. 1–23, Nov. 1973, doi: 10.1146/annurev.es.04.110173.000245.

[35] D. G. Angeler and C. R. Allen, “Quantifying resilience,” *J. Appl. Ecol.*, vol. 53, no. 3, pp. 617–624, 2016.
[36] A. Annarelli and G. Palombi, “Digitalization Capabilities for Sustainable Cyber Resilience: A Conceptual Framework,” *Sustainability*, vol. 13, no. 23, Art. no. 23, Jan. 2021, doi: 10.3390/su132313065.
[37] R. Arghandeh, A. Von Meier, L. Mehrmanesh, and L. Mili, “On the definition of cyber-physical resilience in power systems,” *Renew. Sustain. Energy Rev.*, vol. 58, pp. 1060–1069, May 2016, doi: 10.1016/J.RSER.2015.12.193.
[38] R. M. Bakker, J. Raab, and H. B. Milward, “A preliminary theory of dark network resilience,” *J. Policy Anal. Manage.*, vol. 31, no. 1, pp. 33–62, 2012, doi: 10.1002/pam.20619.
[39] F. Björck, M. Henkel, J. Stirna, and J. Zdravkovic, “Cyber Resilience – Fundamentals for a Definition,” in *New Contributions in Information Systems and Technologies*, A. Rocha, A. M. Correia, S. Costanzo, and L. P. Reis, Eds., Cham: Springer International Publishing, 2015, pp. 311–316. doi: 10.1007/978-3-319-16486-1_31.
[40] E. B. Connelly, C. R. Allen, K. Hatfield, J. M. Palma-Oliveira, D. D. Woods, and I. Linkov, “Features of resilience,” *Environ. Syst. Decis.*, vol. 37, no. 1, pp. 46–50, Mar. 2017, doi: 10.1007/s10669-017-9634-9.
[41] B. Dupont, “The cyber-resilience of financial institutions: significance and applicability,” *J. Cybersecurity*, vol. 5, no. 1, p. tyz013, Jan. 2019, doi: 10.1093/cybsec/tyz013.
[42] B. Dupont, C. Shearing, M. Bernier, and R. Leukfeldt, “The tensions of cyber-resilience: From sensemaking to practice,” *Comput. Secur.*, vol. 132, p. 103372, Sep. 2023, doi: 10.1016/j.cose.2023.103372.
[43] B. D. Fath, C. A. Dean, and H. Katzmair, “Navigating the adaptive cycle: an approach to managing the resilience of social systems,” *Ecol. Soc.*, vol. 20, no. 2, 2015, Accessed: Jul. 29, 2023. [Online]. Available: https://www.jstor.org/stable/26270208
[44] C. Folke, S. R. Carpenter, B. Walker, M. Scheffer, T. Chapin, and J. Rockström, “Resilience Thinking: Integrating Resilience, Adaptability and Transformability,” *Ecol. Soc.*, vol. 15, no. 4, 2010, Accessed: Aug. 26, 2024. [Online]. Available: https://www.jstor.org/stable/26268226
[45] G. Hornor, “Resilience,” *J. Pediatr. Health Care*, vol. 31, no. 3, pp. 384–390, May 2017, doi: 10.1016/j.pedhc.2016.09.005.
[46] D. Jackson, A. Firtko, and M. Edenborough, “Personal resilience as a strategy for surviving and thriving in the face of workplace adversity: a literature review,” *J. Adv. Nurs.*, vol. 60, no. 1, pp. 1–9, 2007, doi: 10.1111/j.1365-2648.2007.04412.x.
[47] C. Johansen, J. Horney, and I. Tien, “Metrics for Evaluating and Improving Community Resilience,” *J. Infrastruct. Syst.*, vol. 23, no. 2, p. 04016032, Jun. 2017, doi: 10.1061/(ASCE)IS.1943-555X.0000329.
[48] A. Kott, M. S. Golan, B. D. Trump, and I. Linkov, “Cyber Resilience: By Design or by Intervention?,” *Computer*, vol. 54, no. 8, pp. 112–117, Aug. 2021, doi: 10.1109/MC.2021.3082836.
[49] A. Kott *et al.*, “Russian Cyber Onslaught was Blunted by Ukrainian Cyber Resilience, not Merely Security,” arXiv.org. Accessed: Sep. 18, 2024. [Online]. Available: https://arxiv.org/abs/2408.14667v1
[50] M. E. Kruk *et al.*, “Building resilient health systems: a proposal for a resilience index,” *BMJ*, p. j2323, May 2017, doi: 10.1136/bmj.j2323.

[51] S. B. Manyena, “The concept of resilience revisited,” *Disasters*, vol. 30, no. 4, pp. 434–450, 2006, doi: 10.1111/j.0361-3666.2006.00331.x.
[52] T. H. Oliver, N. J. B. Isaac, T. A. August, B. A. Woodcock, D. B. Roy, and J. M. Bullock, “Declining resilience of ecosystem functions under biodiversity loss,” *Nat. Commun. 2015 61*, vol. 6, no. 1, pp. 1–8, Dec. 2015, doi: 10.1038/ncomms10122.
[53] R. Patriarca, G. Di Gravio, F. Costantino, A. Falegnami, and F. Bilotta, “An Analytic Framework to Assess Organizational Resilience,” *Saf. Health Work*, vol. 9, no. 3, pp. 265–276, Sep. 2018, doi: 10.1016/J.SHAW.2017.10.005.
[54] F. H. Norris, S. P. Stevens, B. Pfefferbaum, K. F. Wyche, and R. L. Pfefferbaum, “Community Resilience as a Metaphor, Theory, Set of Capacities, and Strategy for Disaster Readiness,” *Am. J. Community Psychol.*, vol. 41, no. 1, pp. 127–150, Mar. 2008, doi: 10.1007/s10464-007-9156-6.
[55] K. A. Pettersen and P. R. Schulman, “Drift, adaptation, resilience and reliability: Toward an empirical clarification,” *Saf. Sci.*, vol. 117, pp. 460–468, Aug. 2019, doi: 10.1016/j.ssci.2016.03.004.
[56] C. Rapaport, T. Hornik-Lurie, O. Cohen, M. Lahad, D. Leykin, and L. Aharonson-Daniel, “The relationship between community type and community resilience,” *Int. J. Disaster Risk Reduct.*, vol. 31, pp. 470–477, Oct. 2018, doi: 10.1016/j.ijdrr.2018.05.020.
[57] R. M. Reischuk, “Cybersecurity: Balancing Efficiency with Long-Term Resilience in Connected Ecosystems,” in *Connected Business*, O. Gassmann and F. Ferrandina, Eds., Cham: Springer International Publishing, 2021, pp. 233–246. doi: 10.1007/978-3-030-76897-3_13.
[58] M. Rutter, “Resilience as a dynamic concept,” *Dev. Psychopathol.*, vol. 24, no. 2, pp. 335–344, May 2012, doi: 10.1017/S0954579412000028.
[59] M. F. Safitra, M. Lubis, and M. T. Kurniawan, “Cyber Resilience: Research Opportunities,” in *Proceedings of the 2023 6th International Conference on Electronics, Communications and Control Engineering*, Fukuoka Japan: ACM, Mar. 2023, pp. 99–104. doi: 10.1145/3592307.3592323.
[60] A. Sharifi, “A critical review of selected tools for assessing community resilience,” *Ecol. Indic.*, vol. 69, pp. 629–647, Oct. 2016, doi: 10.1016/j.ecolind.2016.05.023.
[61] S. M. Southwick, G. A. Bonanno, A. S. Masten, C. Panter-Brick, and R. Yehuda, “Resilience definitions, theory, and challenges: interdisciplinary perspectives,” *Eur. J. Psychotraumatology*, vol. 5, p. 10.3402/ejpt.v5.25338, Oct. 2014, doi: 10.3402/ejpt.v5.25338.
[62] H. T. Tran, M. Balchanos, J. C. Domerçant, and D. N. Mavris, “A framework for the quantitative assessment of performance-based system resilience,” *Reliab. Eng. Syst. Saf.*, vol. 158, pp. 73–84, Feb. 2017, doi: 10.1016/j.ress.2016.10.014.
[63] R. van der Kleij and R. Leukfeldt, “Cyber Resilient Behavior: Integrating Human Behavioral Models and Resilience Engineering Capabilities into Cyber Security,” in *Advances in Human Factors in Cybersecurity*, vol. 960, T. Ahram and W. Karwowski, Eds., in Advances in Intelligent Systems and Computing, vol. 960. , Cham: Springer International Publishing, 2020, pp. 16–27. doi: 10.1007/978-3-030-20488-4_2.
[64] S. E. van der Merwe, R. Biggs, and R. Preiser, “A framework for conceptualizing and assessing the resilience of essential services produced by socio-technical systems,” *Ecol. Soc.*, vol. 23, no. 2, 2018, Accessed: Jan. 10, 2024. [Online]. Available: https://www.jstor.org/stable/26799110

[65] B. Walker, C. S. Holling, S. R. Carpenter, and A. Kinzig, “Resilience, Adaptability and Transformability in Social–ecological Systems,” *Ecol. Soc. Publ. Online Sep 16 2004 Doi105751ES-00650-090205*, vol. 9, no. 2, Sep. 2004, doi: 10.5751/ES-00650-090205.
[66] A. Williams, G. Whiteman, and S. Kennedy, “Cross-Scale Systemic Resilience: Implications for Organization Studies,” *Bus. Soc.*, vol. 60, no. 1, pp. 95–124, Jan. 2021, doi: 10.1177/0007650319825870.
[67] D. D. Woods, “Four concepts for resilience and the implications for the future of resilience engineering,” *Reliab. Eng. Syst. Saf.*, vol. 141, pp. 5–9, Sep. 2015, doi: 10.1016/j.ress.2015.03.018.
[68] X. Zhang, E. Miller-Hooks, and K. Denny, “Assessing the role of network topology in transportation network resilience,” *J. Transp. Geogr.*, vol. 46, pp. 35–45, Jun. 2015, doi: 10.1016/j.jtrangeo.2015.05.006.
[69] D. D. Woods, “The theory of graceful extensibility: basic rules that govern adaptive systems,” *Environ. Syst. Decis.*, vol. 38, no. 4, pp. 433–457, Dec. 2018, doi: 10.1007/s10669-018-9708-3.
[70] D. D. Woods, *Resilience Engineering: Concepts and Precepts*. CRC Press, 2017.
[71] *wordcloud: A little word cloud generator*. Accessed: Nov. 01, 2024. [Online]. Available: https://github.com/amueller/word_cloud
[72] “CountVectorizer,” scikit-learn. Accessed: Nov. 01, 2024. [Online]. Available: https://scikit-learn/stable/modules/generated/sklearn.feature_extraction.text.CountVectorizer.html
[73] S. Niwattanakul, J. Singthongchai, E. Naenudorn, and S. Wanapu, “Using of Jaccard Coefficient for Keywords Similarity,” in *Proceedings of the International MultiConference of Engineers and Computer Scientists*, Hong Kong, China, Mar. 2013. [Online]. Available: https://www.iaeng.org/publication/IMECS2013/IMECS2013_pp380-384.pdf
[74] M. Röder, A. Both, and A. Hinneburg, “Exploring the Space of Topic Coherence Measures,” in *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, in WSDM ’15. New York, NY, USA: Association for Computing Machinery, Feb. 2015, pp. 399–408. doi: 10.1145/2684822.2685324.
[75] J. Chang, S. Gerrish, C. Wang, J. Boyd-graber, and D. Blei, “Reading Tea Leaves: How Humans Interpret Topic Models,” in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2009. Accessed: Mar. 04, 2025. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2009/hash/f92586a25bb3145facd64ab20fd554ff-Abstract.html
[76] B. Fuglede and F. Topsoe, “Jensen-Shannon divergence and Hilbert space embedding,” in *International Symposium onInformation Theory, 2004. ISIT 2004. Proceedings.*, Jun. 2004, pp. 31-. doi: 10.1109/ISIT.2004.1365067.
[77] J. H. Holland, “Complex Adaptive Systems,” *Daedalus*, vol. 121, no. 1, pp. 17–30, 1992.
[78] I. Linkov and A. Kott, “Fundamental Concepts of Cyber Resilience: Introduction and Overview,” in *Cyber Resilience of Systems and Networks*, A. Kott and I. Linkov, Eds., Cham: Springer International Publishing, 2019, pp. 1–25. doi: 10.1007/978-3-319-77492-3_1.
[79] “Review of the Summer 2023 Microsoft Exchange Online Intrusion,” Cybersecurity & Infrastructure Security Agency, 2023. [Online]. Available:

https://www.cisa.gov/sites/default/files/2024-04/CSRB_Review_of_the_Summer_2023_MEO_Intrusion_Final_508c.pdf
[80] A. Ribeiro, “Trump administration dismantles CSRB, leaves future of cybersecurity oversight in question,” Industrial Cyber. Accessed: Mar. 04, 2025. [Online]. Available: https://industrialcyber.co/regulation-standards-and-compliance/trump-administration-dismantles-csrb-leaves-future-of-cybersecurity-oversight-in-question/
[81] D. P. Aldrich and M. A. Meyer, “Social Capital and Community Resilience,” *Am. Behav. Sci.*, vol. 59, no. 2, pp. 254–269, Feb. 2015, doi: 10.1177/0002764214550299.
[82] N. Leveson, N. Dulac, K. Marais, and J. Carroll, “Moving Beyond Normal Accidents and High Reliability Organizations: A Systems Approach to Safety in Complex Systems,” *Organ. Stud.*, vol. 30, no. 2–3, pp. 227–249, Feb. 2009, doi: 10.1177/0170840608101478.
[83] K. E. Weick, “Organizational Culture as a Source of High Reliability,” *Calif. Manage. Rev.*, vol. 29, no. 2, pp. 112–127, Jan. 1987, doi: 10.2307/41165243.
[84] J. D. Orton and K. E. Weick, “Loosely Coupled Systems: A Reconceptualization,” *Acad. Manage. Rev.*, vol. 15, no. 2, pp. 203–223, Apr. 1990, doi: 10.5465/amr.1990.4308154.
[85] C. Nemeth and R. Cook, “Reliability versus Resilience: What Does Healthcare Need?,” *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 51, no. 11, pp. 621–625, Oct. 2007, doi: 10.1177/154193120705101104.
[86] C. P. Nemeth and E. Hollnagel, Eds., *Advancing Resilient Performance*. Cham: Springer International Publishing, 2022. doi: 10.1007/978-3-030-74689-6.
[87] T. K. Haavik, S. Antonsen, R. Rosness, and A. Hale, “HRO and RE: A pragmatic perspective,” *Saf. Sci.*, vol. 117, pp. 479–489, Aug. 2019, doi: 10.1016/j.ssci.2016.08.010.
[88] C. J. Foster, K. L. Plant, and N. A. Stanton, “Adaptation as a source of safety in complex socio-technical systems: A literature review and model development,” *Saf. Sci.*, vol. 118, pp. 617–631, Oct. 2019, doi: 10.1016/j.ssci.2019.05.035.
[89] A. S. Downing, E. H. van Nes, W. M. Mooij, and M. Scheffer, “The Resilience and Resistance of an Ecosystem to a Collapse of Diversity,” *PLOS ONE*, vol. 7, no. 9, p. e46135, Sep. 2012, doi: 10.1371/JOURNAL.PONE.0046135.
[90] B. Walker, A. Kinzig, and J. Langridge, “Plant Attribute Diversity, Resilience, and Ecosystem Function: The Nature and Significance of Dominant and Minor Species”.
[91] J. J. Stachowicz and J. E. Byrnes, “Species diversity, invasion success, and ecosystem functioning: Disentangling the influence of resource competition, facilitation, and extrinsic factors,” *Mar. Ecol. Prog. Ser.*, vol. 311, pp. 251–262, 2006, doi: 10.3354/meps311251.
[92] D. J. Rapport, W. G. Whitford, and M. Hildén, “Common Patterns of Ecosystem Breakdown Under Stress,” *Environ. Monit. Assess.*, vol. 51, no. 1/2, pp. 171–178, 1998, doi: 10.1023/A:1005935202518.
[93] T. H. Oliver *et al.*, “Biodiversity and Resilience of Ecosystem Functions,” *Trends Ecol. Evol.*, vol. 30, no. 11, pp. 673–684, Nov. 2015, doi: 10.1016/j.tree.2015.08.009.
[94] D. J. Rapport, H. A. Regier, and T. C. Hutchinson, “Ecosystem Behavior Under Stress,” *Am. Nat.*, vol. 125, no. 5, pp. 617–640, May 1985, doi: 10.1086/284368.
[95] I. M. Côté and E. S. Darling, “Rethinking Ecosystem Resilience in the Face of Climate Change,” *PLOS Biol.*, vol. 8, no. 7, p. e1000438, Jul. 2010, doi: 10.1371/journal.pbio.1000438.
[96] E. Hollnagel, D. Woods, and N. Leveson, “Resilience Engineering : Concepts and Precepts,” *Resil. Eng. Concepts Precepts*, Feb. 2006.

[97] C. R. Allen, D. G. Angeler, A. S. Garmestani, L. H. Gunderson, and C. S. Holling, "Panarchy: Theory and Application," *Ecosystems*, vol. 17, no. 4, pp. 578–589, Jun. 2014, doi: 10.1007/s10021-013-9744-2.
[98] C. S. Holling, "Surprise for Science, Resilience for Ecosystems, and Incentives for People," *Ecol. Appl.*, vol. 6, no. 3, pp. 733–735, Aug. 1996, doi: 10.2307/2269475.
[99] C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems*, vol. 4, no. 5, pp. 390–405, Aug. 2001, doi: 10.1007/s10021-001-0101-5.
[100] C. Folke, J. Colding, and F. Berkes, "Building resilience and adaptive capacity in social-ecological systems".
[101] E. Geller, "The US Government Has a Microsoft Problem," *Wired*. Accessed: Jun. 19, 2024. [Online]. Available: https://www.wired.com/story/the-us-government-has-a-microsoft-problem/
[102] Z. Siddiqui, "US State Dept broadens security vendor list amid Microsoft hacking woes," *Reuters*, May 07, 2024. Accessed: Jun. 19, 2024. [Online]. Available: https://www.reuters.com/technology/cybersecurity/us-state-dept-broadens-security-vendor-list-amid-microsoft-hacking-woes-2024-05-07/
[103] E. Kovacs, "Microsoft Criticized Over Handling of Critical Power Platform Vulnerability," SecurityWeek. Accessed: Jun. 19, 2024. [Online]. Available: https://www.securityweek.com/microsoft-criticized-over-handling-of-critical-power-platform-vulnerability/
[104] S. Aren and H. Nayman Hamamcı, "Biases in managerial decision making: Regret aversion, endowment, confirmation, self-control, recency," vol. 8, pp. 62–69, Jul. 2021.
[105] H.-M. Darley, "Dangers of succumbing to bias in cyber security : An evaluation of the impact of cognitive biases on threat assessments and cyber security strategies," *Cyber Secur. Peer-Rev. J.*, vol. 6, no. 3, pp. 211–219, Jan. 2023.
[106] M. Schwartz, *A Seat at the Table*. IT Revolution, 2017. Accessed: May 06, 2024. [Online]. Available: https://itrevolution.com/product/a-seat-at-the-table/
[107] M. X. Heiligenstein, "Microsoft Data Breaches: Full Timeline Through 2024," Firewall Times. Accessed: Jun. 19, 2024. [Online]. Available: https://firewalltimes.com/microsoft-data-breach-timeline/
[108] R. Naraine, "Microsoft's Security Chickens Have Come Home to Roost," SecurityWeek. Accessed: Jun. 19, 2024. [Online]. Available: https://www.securityweek.com/microsofts-security-chickens-have-come-home-to-roost/
[109] B. Levin and L. Downes, "Microsoft, Google, and a New Era of Antitrust," *Harvard Business Review*, Feb. 17, 2023. Accessed: Jun. 19, 2024. [Online]. Available: https://hbr.org/2023/02/microsoft-google-and-a-new-era-of-antitrust
[110] R. D. Burke Doris, "Microsoft Chose Profit Over Security and Left U.S. Government Vulnerable to Russian Hack, Whistleblower Says," ProPublica. Accessed: Jun. 14, 2024. [Online]. Available: https://www.propublica.org/article/microsoft-solarwinds-golden-saml-data-breach-russian-hackers
[111] J. Collins, *Good to Great: Why Some Companies Make the Leap...And Others Don't*, First Edition. New York, NY: Harper Business, 2001.
[112] S. J. Spear, *The High-Velocity Edge: How Market Leaders Leverage Operational Excellence to Beat the Competition*, 2nd edition. McGraw Hill, 2010.

[113] E. Terrell, “Finding General Electric | Inside Adams,” The Library of Congress. Accessed: May 16, 2024. [Online]. Available: https://blogs.loc.gov/inside_adams/2019/09/ge
[114] “About the Microsoft Security Development Lifecycle.” Accessed: Jun. 21, 2024. [Online]. Available: https://www.microsoft.com/en-us/securityengineering/sdl/about
[115] “Memo from Bill Gates,” Stories. Accessed: Apr. 29, 2024. [Online]. Available: https://news.microsoft.com/2012/01/11/memo-from-bill-gates/
[116] B. Smith, “A new world of security: Microsoft’s Secure Future Initiative,” Microsoft On the Issues. Accessed: Apr. 29, 2024. [Online]. Available: https://blogs.microsoft.com/on-the-issues/2023/11/02/secure-future-initiative-sfi-cybersecurity-cyberattacks/
[117] R. Cook and J. Rasmussen, “‘Going solid’: a model of system dynamics and consequences for patient safety,” *BMJ Qual. Saf.*, vol. 14, no. 2, pp. 130–134, Apr. 2005, doi: 10.1136/qshc.2003.009530.
[118] “Glossary of Security Terms | SANS Institute.” Accessed: Jul. 11, 2024. [Online]. Available: https://www.sans.org/security-resources/glossary-of-terms/
[119] C. S. E. Canada, “Glossary,” Canadian Centre for Cyber Security. Accessed: Jul. 11, 2024. [Online]. Available: https://www.cyber.gc.ca/en/glossary
[120] “Vocabulary | NICCS.” Accessed: Jul. 11, 2024. [Online]. Available: https://niccs.cisa.gov/cybersecurity-career-resources/vocabulary
[121] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter,” Mar. 01, 2020, *arXiv*: arXiv:1910.01108. doi: 10.48550/arXiv.1910.01108.
[122] E. Aghaei, X. Niu, W. Shadid, and E. Al-Shaer, “SecureBERT: A Domain-Specific Language Model for Cybersecurity,” in *Security and Privacy in Communication Networks*, F. Li, K. Liang, Z. Lin, and S. K. Katsikas, Eds., Cham: Springer Nature Switzerland, 2023, pp. 39–56. doi: 10.1007/978-3-031-25538-0_3.
[123] “distilbert/distilbert-base-uncased-finetuned-sst-2-english · Hugging Face.” Accessed: Jan. 18, 2025. [Online]. Available: https://huggingface.co/distilbert/distilbert-base-uncased-finetuned-sst-2-english
[124] “ehsanaghaei/SecureBERT_Plus · Hugging Face.” Accessed: Jan. 18, 2025. [Online]. Available: https://huggingface.co/ehsanaghaei/SecureBERT_Plus
[125] “json — JSON encoder and decoder,” Python documentation. Accessed: Jan. 19, 2025. [Online]. Available: https://docs.python.org/3/library/json.html
[126] *matplotlib: Python plotting package*. Python. Accessed: Jan. 19, 2025. [Online]. Available: https://matplotlib.org
[127] *seaborn: Statistical data visualization*. Python.
[128] *transformers: State-of-the-art Machine Learning for JAX, PyTorch and TensorFlow*. Python. Accessed: Jan. 19, 2025. [OS Independent]. Available: https://github.com/huggingface/transformers
[129] *pandas: Powerful data structures for data analysis, time series, and statistics*. Cython, Python. Accessed: Jan. 19, 2025. [OS Independent]. Available: https://pandas.pydata.org
[130] “NumPy.” Accessed: Jan. 19, 2025. [Online]. Available: https://numpy.org/
[131] *scikit-learn: A set of python modules for machine learning and data mining*. C, Python.
[132] “Claude 3.5 Sonnet,” Anthropic. Accessed: Jan. 20, 2025. [Online]. Available: https://claude.ai/new

[133] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks,” Aug. 27, 2019, *arXiv*: arXiv:1908.10084. doi: 10.48550/arXiv.1908.10084.
[134] “SentenceTransformers Documentation — Sentence Transformers documentation.” Accessed: Jan. 20, 2025. [Online]. Available: https://sbert.net/
[135] 14:00-17:00, “ISO 22301:2019,” ISO. Accessed: Jul. 06, 2024. [Online]. Available: https://www.iso.org/standard/75106.html
[136] 14:00-17:00, “ISO/IEC 27001:2022,” ISO. Accessed: Jul. 06, 2024. [Online]. Available: https://www.iso.org/standard/27001
[137] 14:00-17:00, “ISO/IEC 27002:2022,” ISO. Accessed: Jul. 06, 2024. [Online]. Available: https://www.iso.org/standard/75652.html
[138] “ISO 31000:2018,” ISO. Accessed: Jan. 18, 2025. [Online]. Available: https://www.iso.org/standard/65694.html
[139] National Institute of Standards and Technology, “The NIST Cybersecurity Framework (CSF) 2.0,” National Institute of Standards and Technology, Gaithersburg, MD, NIST CSWP 29, Feb. 2024. doi: 10.6028/NIST.CSWP.29.
[140] “Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-37 Rev. 2, Dec. 2018. doi: 10.6028/NIST.SP.800-37r2.
[141] M. Merritt, S. Hansche, B. Ellis, K. Sanchez-Cherry, J. Snyder, and D. Walden, “Building a Cybersecurity and Privacy Learning Program,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-50 Rev. 1 (Draft), Aug. 2023. doi: 10.6028/NIST.SP.800-50r1.ipd.
[142] A. Nelson, S. Rekhi, M. Souppaya, and K. Scarfone, “Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-61 Rev. 3 (Draft), Apr. 2024. doi: 10.6028/NIST.SP.800-61r3.ipd.
[143] K. Stouffer *et al.*, “Guide to Operational Technology (OT) Security,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-82 Rev. 3, Sep. 2023. doi: 10.6028/NIST.SP.800-82r3.
[144] R. Ross, M. McEvilley, and J. Carrier Oren, “Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems,” National Institute of Standards and Technology, NIST SP 800-160, Nov. 2016. doi: 10.6028/NIST.SP.800-160.
[145] J. Boyens, A. Smith, N. Bartol, K. Winkler, A. Holbrook, and M. Fallon, “Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-161 Rev. 1, May 2022. doi: 10.6028/NIST.SP.800-161r1.
[146] A. Regenscheid, “Platform Firmware Resiliency Guidelines,” National Institute of Standards and Technology, NIST Special Publication (SP) 800-193, May 2018. doi: 10.6028/NIST.SP.800-193.
[147] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” National Institute of Standards and Technology, Aug. 2020. doi: 10.6028/NIST.SP.800-207.
[148] S. Quinn *et al.*, “Information and Communications Technology (ICT) Risk Outcomes: Integrating ICT Risk Management Programs with the Enterprise Risk Portfolio,” National

Institute of Standards and Technology, NIST Special Publication (SP) 800-221A, Nov. 2023. doi: 10.6028/NIST.SP.800-221A.
[149] "Implementing a Zero Trust Architecture," National Institute of Standards and Technology, NIST Special Publication (SP) 1800-35 (Draft), Aug. 2023. Accessed: Jul. 06, 2024. [Online]. Available: https://csrc.nist.gov/pubs/sp/1800/35/3prd
[150] "Cyber Assessment Framework." National Cyber Security Centre, Apr. 15, 2024. Accessed: Jul. 06, 2024. [Online]. Available: https://www.ncsc.gov.uk/collection/cyber-assessment-framework/introduction-to-caf
[151] *BSI-Standard 100-1: Information Security Management Systems (ISMS)*. Accessed: Jul. 06, 2024. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/BSIStandards/standard_100-1_e_pdf.html?nn=132646
[152] *BSI-Standard 200-1 - Information Security Management Systems (ISMS)*.
[153] *BSI-Standard 200-2 - IT-Grundschutz Methodology*.
[154] *BSI-Standard 200-3 - Risk Analysis based on IT-Grundschutz*.
[155] "IT-Grundschutz-Compendium." German Federal Office for Information Security, Feb. 2022. Accessed: Jul. 06, 2024. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Grundschutz/International/bsi_it_gs_comp_2022.pdf?__blob=publicationFile&v=2
[156] "CIS Critical Security Controls Version 8.1," CIS. Accessed: Jul. 06, 2024. [Online]. Available: https://www.cisecurity.org/controls/
[157] "Cybersecurity Maturity Model Certification." Accessed: Jul. 06, 2024. [Online]. Available: https://dodcio.defense.gov/CMMC/
[158] "COBIT 5." 2019. Accessed: Jul. 06, 2024. [Online]. Available: https://store.isaca.org/s/store#/store/browse/detail/a2S4w000004KoCDEA0
[159] "ITIL 4: the framework for the management of IT-enabled services." Axelos, 2019. Accessed: Jul. 06, 2024. [Online]. Available: https://www.axelos.com/certifications/itil-service-management
[160] "Energy Sector Cybersecurity Framework Implementation Guidance," U.S. Department of Energy, Jan. 2015. Accessed: Jan. 18, 2025. [Online]. Available: https://www.energy.gov/ceser/articles/energy-sector-cybersecurity-framework-implementation-guidance
[161] "FFIEC Cybersecurity Assessment Tool." 2017. Accessed: Jul. 06, 2024. [Online]. Available: https://www.ffiec.gov/cyberassessmenttool.htm#tool
[162] "OCTAVE Criteria, Version 2.0." Accessed: Jul. 06, 2024. [Online]. Available: https://insights.sei.cmu.edu/library/octave-criteria-version-20/
[163] T. W. House, "Executive Order on Strengthening and Promoting Innovation in the Nation's Cybersecurity," The White House. Accessed: Jan. 18, 2025. [Online]. Available: https://www.whitehouse.gov/briefing-room/presidential-actions/2025/01/16/executive-order-on-strengthening-and-promoting-innovation-in-the-nations-cybersecurity/
[164] "The Cyber Resilience Index: Advancing Organizational Cyber Resilience," World Economic Forum. Accessed: Jul. 06, 2024. [Online]. Available: https://www.weforum.org/publications/the-cyber-resilience-index-advancing-organizational-cyber-resilience/
[165] D. Bodeau, R. Graubart, W. Heinbockel, and E. Laderman, "Cyber Resiliency Engineering Aid—The Updated Cyber Resiliency Engineering Framework and Guidance

on Applying Cyber Resiliency Techniques," May 2015, Accessed: Jul. 06, 2024. [Online]. Available: https://www.mitre.org/news-insights/publication/cyber-resiliency-engineering-aid-updated-cyber-resiliency-engineering
[166] E. Laderman, D. Bodeau, R. Graubart, and L. K. Jones, "Cyber Resiliency Framework and Cyber Survivability Attributes," Jan. 2024, Accessed: Jul. 06, 2024. [Online]. Available: https://www.mitre.org/news-insights/publication/cyber-resiliency-framework-and-cyber-survivability-attributes
[167] "Cyber Resiliency Level®," Lockheed Martin. Accessed: Jul. 07, 2024. [Online]. Available: https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-resiliency-level.html
[168] "Operational Resilience Framework," Global Resilience Foundation. Accessed: May 18, 2024. [Online]. Available: https://www.grf.org/orf
[169] "Cybersecurity Culture Guidelines: Behavioural Aspects of Cybersecurity," Report/Study. Accessed: Jul. 06, 2024. [Online]. Available: https://www.enisa.europa.eu/publications/cybersecurity-culture-guidelines-behavioural-aspects-of-cybersecurity
[170] "National Resilience Strategy," Jan. 2025. Accessed: Jan. 19, 2025. [Online]. Available: https://www.whitehouse.gov/briefing-room/statements-releases/2025/01/18/national-resilience-strategy/
[171] O. Reznikova, "National Resilience In a Changing Security Environment," National Institute for Strategic Studies, Kyiv, Ukraine. [Online]. Available: https://www.marshallcenter.org/sites/default/files/files/2023-01/National%20Resilience_EN.pdf
[172] "National Climate Resilience Framework," Sep. 2023. Accessed: Jan. 18, 2025. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2023/09/National-Climate-Resilience-Framework-FINAL.pdf
[173] "Health Emergency and Disaster Risk Management Framework." Accessed: Jan. 18, 2025. [Online]. Available: https://www.who.int/publications/i/item/9789241516181
[174] "Principles for operational resilience," Mar. 2021. Accessed: Jan. 18, 2025. [Online]. Available: https://www.bis.org/bcbs/publ/d516.htm
[175] "Assessing Resilinece in Social-Ecological Systems: Workbook for Practitioners," 2010. Accessed: Jan. 18, 2025. [Online]. Available: https://www.resalliance.org/files/ResilienceAssessmentV2_2.pdf
[176] "Community Resilience System Initiative Steering Committee - Final Report," Community & Regional Resilience Institute, Aug. 2011. [Online]. Available: https://merid.org/wp-content/uploads/2019/08/CRSI-Final-Report.pdf
[177] "APEC Disaster Risk Reduction Framework," Nov. 2015. Accessed: Jan. 18, 2025. [Online]. Available: https://www.apec.org/docs/default-source/groups/epwg/2024/apecdisasterriskreductionframework_endorsed.pdf?sfvrsn=8d9ce067_2
[178] *Reliability Standards for the Bulk Electric Systems of North America*, Jun. 03, 2024. Accessed: Jul. 06, 2024. [Online]. Available: https://www.nerc.com/pa/Stand/Pages/default.aspx
[179] "Infrastructure Resilience Planning Framework (IRPF)," 2024.
[180] "National Infrastructure Protection Plan," Cybersecurity & Infrastructure Security Agency, 2013. Accessed: Jan. 18, 2025. [Online]. Available:

https://www.cisa.gov/sites/default/files/2022-11/national-infrastructure-protection-plan-2013-508.pdf
[181] “Guidelines for Resilience Systems Analysis: How to Analyse Risk and Build a Roadmap to Resilience,” OECD. Accessed: Jan. 18, 2025. [Online]. Available: https://www.oecd.org/en/publications/guidelines-for-resilience-systems-analysis-how-to-analyse-risk-and-build-a-roadmap-to-resilience_3b1d3efe-en.html
[182] “glob — Unix style pathname pattern expansion,” Python documentation. Accessed: Jan. 26, 2025. [Online]. Available: https://docs.python.org/3/library/glob.html
[183] D. W. Hubbard and R. Seiersen, *How to Measure Anything in Cybersecurity Risk*. John Wiley & Sons, 2016.
[184] “Critical Infrastructure Sectors.” Accessed: Jun. 20, 2024. [Online]. Available: https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors

APPENDIX 1: FINAL DICTIONARY AND CATEGORIZATION OF CYBER RESILIENCE TERMINOLOGY

Custom Stop Word Dictionary

The custom stop word dictionary was developed iteratively during tf*idf analyses. It contains: mark, color, annotation, library, pdf, yellow, highlight, note, highlighting, highlighted, highlighter, highlighters, highlighted, highlighting, item, one, two, three, four, five, six, seven, eight, nine, ten, eleven, twelve, thirteen, fourteen, fifteen, sixteen, seventeen, eighteen, nineteen, twenty, may, approach, also, based, author, different, note, event, used, new, way, end, term, within, use, note, many, however, set, high, well, begin, calendar, generator, voltage, relay, shed, whichever, swing, equal, and impedance.

Final Time-Scale Classification Dictionary

Term	Time Classifier	Scale Classifier
aad	Immediate	Technical
acceptable use agreement	Immediate	Sociotechnical
access	Short	Sociotechnical
access control mechanism	Short	Sociotechnical
access level	Short	Sociotechnical
access type	Short	Sociotechnical
access vector	Short	Sociotechnical
account	Immediate	Sociotechnical
accredit	Immediate	Technical
accuracy	Immediate	Technical
accuracy (absolute)	Immediate	Technical
accuracy (relative)	Immediate	Technical
acm	Immediate	Technical
active attack	Immediate	Technical
active content	Immediate	Technical
active state	Immediate	Technical

active tag	Immediate	Technical
activity	Immediate	Technical
actuating capability	Long	Technical
adaptability	Immediate	Technical
adaptive capacity	Mid	Technical
adaptive cycle	Mid	Organizational
adequate security	Immediate	Technical
adjudicative entity	Immediate	Technical
advanced cyber threat	Short	Sociotechnical
advanced technology attachment	Mid	Technical
adversary	Immediate	Technical
adverse consequence	Immediate	Technical
ae	Immediate	Technical
aes	Immediate	Technical
agency dashboard	Immediate	Sociotechnical
agility	Immediate	Technical
agreement	Immediate	Technical
air gap	Immediate	Technical
algorithm	Immediate	Technical
all-source intelligence	Mid	Sociotechnical
ambiguity rule	Immediate	Technical
analysis	Short	Technical
anonymized information	Immediate	Community/Sector
anonymizers	Immediate	Technical
anti-forensic	Immediate	Technical
anti-spoof	Immediate	Technical
Anti-virus software	Immediate	Technical
aperiodic templates test	Immediate	Technical
applicability statement	Immediate	Technical
application translation	Immediate	Technical
application-proxy gateway	Immediate	Technical
approval status	Immediate	Technical
approval to operate	Immediate	Technical
approved	Immediate	Technical
apt	Immediate	Technical
architecture constructs	Mid	Technical
architecture description	Mid	Sociotechnical
architecture design principles	Mid	Sociotechnical
architecture framework	Mid	Sociotechnical
archive	Immediate	Technical
Artificial intelligence	Immediate	Technical
assembly	Immediate	Technical
assessment	Mid	Technical

assessment criterion/criteria	Mid	Technical
assessment element attribute	Mid	Technical
asset	Immediate	Technical
asset identification	Immediate	Technical
asset reporting format	Immediate	Technical
asset tag	Immediate	Technical
assignment statement	Immediate	Technical
associated data	Immediate	Technical
assurance	Mid	Technical
assurance case	Mid	Technical
assurance evidence	Mid	Organizational
assurance message	Mid	Technical
assurance of domain parameter validity	Mid	Technical
assurance of integrity	Mid	Technical
assurance of possession	Mid	Technical
assurance of public key validity	Mid	Technical
assurance of validity	Mid	Technical
assurance- signature	Immediate	Technical
assured information sharing	Mid	Community/Sector
assured software	Immediate	Technical
ata	Immediate	Technical
attack	Immediate	Technical
attack method	Immediate	Technical
attack pattern	Short	Sociotechnical
attack signature	Immediate	Technical
attack tree	Immediate	Technical
attacker	Immediate	Technical
attribute	Immediate	Technical
attribute authority	Immediate	Organizational
attribute bundle	Immediate	Technical
attribute-value pair	Immediate	Technical
audience	Immediate	Technical
audit administrator	Immediate	Sociotechnical
audit log	Immediate	Technical
audit record	Immediate	Technical
audit reduction tools	Immediate	Technical
auditor	Short	Organizational
authenticable entity	Immediate	Technical
authenticate	Immediate	Technical
authenticated decryption	Immediate	Technical
authenticated encryption	Immediate	Technical
authentication mechanism	Immediate	Technical
authenticity	Immediate	Technical

author	Short	Sociotechnical
authoritative source	Immediate	Technical
authority to operate	Immediate	Organizational
authorized	Immediate	Sociotechnical
automated security monitoring	Immediate	Technical
availability	Immediate	Technical
availability impact	Immediate	Technical
avp	Immediate	Technical
awareness	Immediate	Technical
back-channel communication	Immediate	Sociotechnical
backdoor	Immediate	Technical
bad	Immediate	Technical
banner grabbing	Immediate	Technical
base point	Immediate	Technical
base standards	Long	Organizational
baseline security	Immediate	Technical
basis vector	Immediate	Technical
bastion host	Immediate	Technical
beacon	Immediate	Technical
Beaconing	Immediate	Technical
benchmark producer	Immediate	Technical
benign environment	Immediate	Technical
ber-tlv data object	Immediate	Technical
best practice	Short	Community/Sector
bias	Immediate	Technical
biased	Immediate	Technical
bi-directional (cds)	Immediate	Technical
binary sequence	Immediate	Technical
binding	Immediate	Technical
binomial distribution	Immediate	Technical
bio	Immediate	Technical
biometric authentication (bio, bio-a)	Immediate	Technical
bit	Immediate	Technical
bit error	Immediate	Technical
bit error rate	Immediate	Technical
bit length	Immediate	Technical
bit stream imaging	Immediate	Technical
bit string	Immediate	Technical
black	Immediate	Technical
black data	Immediate	Technical
blacklist	Immediate	Technical
blacklisting	Immediate	Sociotechnical
block	Immediate	Technical

block cipher	Immediate	Technical
block cipher algorithm	Immediate	Technical
block data	Immediate	Technical
block header	Immediate	Technical
blockchain implementation	Short	Technical
blockchain network	Immediate	Technical
blockchain technology	Mid	Technical
blocklist	Immediate	Technical
body of evidence	Immediate	Technical
botnet	Immediate	Technical
Boundary interface	Immediate	Technical
boundary protection	Immediate	Technical
breach	Immediate	Technical
breadth	Immediate	Technical
breakdown structure	Immediate	Technical
broad network access	Immediate	Technical
brokered trust	Immediate	Technical
Browser-based exploitation	Immediate	Sociotechnical
buffer overflow attack	Immediate	Technical
bug	Short	Technical
bug bounty	Short	Technical
build security in	Immediate	Technical
buyer	Mid	Organizational
byte	Immediate	Technical
call back	Immediate	Technical
canister (comsec)	Immediate	Technical
capabilities catalog	Immediate	Technical
capability	Long	Technical
capability list	Long	Technical
capability, behavior management	Long	Sociotechnical
capability, boundary management	Long	Sociotechnical
capability, event preparation management	Long	Sociotechnical
capability, hardware asset management	Long	Technical
capability, iscm	Immediate	Technical
capability, manage and assess risk	Mid	Organizational
capability, perform resilient systems engineering	Generational	Technical
capability, privilege and account management	Long	Sociotechnical
capability, security	Immediate	Technical
capability, software asset management	Long	Technical
capability, trust management	Long	Sociotechnical

capture	Immediate	Technical
card verifiable	Immediate	Technical
catalog	Immediate	Technical
category	Immediate	Technical
cbeff sub-header	Immediate	Technical
ccb	Immediate	Technical
cdm	Immediate	Technical
cd-rewritable	Immediate	Technical
cd-rw	Immediate	Technical
central office of record	Immediate	Technical
central oversight authority	Immediate	Organizational
central services node	Immediate	Technical
certificate	Immediate	Technical
certificate management	Long	Sociotechnical
certificate transparency	Immediate	Technical
certificate-inventory management	Long	Sociotechnical
certificate-related information	Immediate	Community/Sector
certification	Long	Technical
certification analyst	Long	Technical
certifier	Immediate	Technical
chain	Immediate	Technical
chain of custody	Immediate	Technical
chain of trust	Immediate	Technical
chaining	Immediate	Technical
challenge-response protocol	Immediate	Technical
check word	Immediate	Technical
checking disabled	Immediate	Technical
checklist	Immediate	Technical
checklist type	Immediate	Technical
checksum	Immediate	Technical
choreography	Immediate	Sociotechnical
cio	Immediate	Technical
ckms designer	Immediate	Technical
ckms developer	Short	Sociotechnical
ckms hierarchy	Immediate	Technical
claimed address	Immediate	Technical
clean host	Immediate	Technical
clean word list	Immediate	Technical
clear	Immediate	Technical
cleartext	Immediate	Technical
client application	Immediate	Technical
client node	Immediate	Technical
clock	Immediate	Technical

closed source operating system	Immediate	Technical
closed storage	Immediate	Technical
closed system	Immediate	Technical
cluster	Immediate	Technical
cmac	Immediate	Technical
cmrr	Immediate	Technical
code	Short	Technical
code vocabulary	Short	Technical
codec	Immediate	Technical
cold site	Short	Organizational
collection system	Immediate	Technical
collision	Immediate	Technical
collision resistance	Immediate	Technical
commercial-off-the-shelf (cots)	Immediate	Technical
committee draft	Immediate	Technical
commodity service	Immediate	Technical
common control provider	Immediate	Technical
common platform enumeration (cpe)	Immediate	Technical
common services provider (csp)	Immediate	Sociotechnical
common vulnerabilities and exposures identifiers	Immediate	Technical
communications deception	Immediate	Technical
communications router	Immediate	Technical
community of interest (coi)	Immediate	Community/Sector
compensating security control	Immediate	Technical
competency	Immediate	Technical
competent security official	Immediate	Technical
component specification	Immediate	Technical
compromise	Immediate	Technical
compromise recovery	Short	Technical
compromised key list (ckl)	Immediate	Technical
compromised state	Immediate	Technical
Compromising emanations	Immediate	Technical
computer abuse	Immediate	Technical
computer network attack (cna)	Immediate	Sociotechnical
computer network exploitation (cne)	Immediate	Sociotechnical
computer network operations (cno)	Short	Sociotechnical
computer security subsystem	Immediate	Technical
COMSEC	Immediate	Technical
comsec account manager	Short	Sociotechnical
COMSEC incident	Short	Sociotechnical
comsec insecurity	Immediate	Technical
COMSEC material	Immediate	Technical

comsec monitoring	Immediate	Sociotechnical
comsec profile	Immediate	Technical
comsec system data	Immediate	Technical
concept of operations	Mid	Sociotechnical
concept relationship style	Immediate	Technical
condition coverage	Immediate	Technical
confidentiality impact	Immediate	Technical
configurable	Immediate	Technical
configuration baseline	Short	Technical
configuration settings	Short	Technical
conflict	Immediate	Technical
conflict resolution	Short	Technical
confluent hypergeometric function	Immediate	Technical
consensus model	Immediate	Technical
consent banner	Immediate	Technical
consequence	Immediate	Technical
consortium	Long	Community/Sector
container runtime	Immediate	Technical
contamination	Immediate	Technical
content signing certificate	Immediate	Technical
content type	Immediate	Technical
continuity of operations plan	Immediate	Sociotechnical
continuous monitoring as a service	Immediate	Technical
control algorithm	Immediate	Technical
control baseline	Immediate	Technical
control effectiveness	Immediate	Technical
control enhancement	Immediate	Technical
controlled cryptographic item (cci) assembly	Immediate	Technical
controlled cryptographic item (cci) component	Immediate	Technical
controlled space	Immediate	Technical
coordination	Immediate	Technical
copy (data)	Immediate	Technical
cor	Immediate	Technical
corporate-owned personally-enabled (cope)	Immediate	Technical
correct re-identifications	Immediate	Technical
correctness proof	Immediate	Technical
counterfeit	Immediate	Technical
counterintelligence	Long	Sociotechnical
countermeasures	Immediate	Technical
course of action	Immediate	Technical

covert channel	Immediate	Technical
covert storage channel	Immediate	Technical
cp	Immediate	Technical
critical	Immediate	Technical
critical asset	Immediate	Technical
critical value	Immediate	Technical
crl	Immediate	Technical
cross domain baseline list	Immediate	Technical
cross domain sunset list	Immediate	Technical
cryptanalysis	Short	Sociotechnical
cryptographic algorithm	Immediate	Technical
cryptographic api: next generation	Immediate	Technical
cryptographic application	Immediate	Technical
cryptographic checksum	Immediate	Technical
cryptographic ignition key	Immediate	Technical
cryptographic key	Immediate	Technical
cryptographic randomization	Immediate	Technical
cryptographic system review	Immediate	Technical
csf category	Immediate	Technical
csf core	Immediate	Technical
csf organizational profile	Long	Organizational
csm	Immediate	Technical
csn	Immediate	Technical
csp	Immediate	Technical
csr	Immediate	Technical
cut	Immediate	Technical
cve	Immediate	Technical
Cyber attack	Immediate	Technical
cyber attack	Immediate	Technical
cyber incident	Short	Sociotechnical
cyber resiliency construct	Mid	Organizational
cyber resiliency control	Immediate	Technical
Cyber security	Immediate	Technical
Cyberattack	Immediate	Technical
cybersecurity event	Immediate	Technical
cybersecurity risks throughout the supply chain	Immediate	Community/Sector
cybersecurity supply chain risk assessment	Mid	Organizational
cybersecurity-aware	Short	Sociotechnical
dac	Immediate	Technical
damage	Immediate	Technical
dao	Immediate	Technical

dashboard	Immediate	Technical
data asset	Immediate	Technical
data breach	Immediate	Organizational
data collector	Immediate	Technical
data encryption standard	Immediate	Technical
data link layer	Immediate	Technical
data loss	Immediate	Technical
data mining	Immediate	Technical
data origin authentication	Immediate	Technical
data provenance	Immediate	Technical
data spillage	Immediate	Sociotechnical
data theft	Immediate	Sociotechnical
data transfer solution	Immediate	Technical
data universe	Immediate	Technical
data-encryption key	Immediate	Technical
db	Immediate	Technical
DBaaS	Immediate	Technical
DDoS	Immediate	Technical
decentralized autonomous organization	Immediate	Technical
decertification	Immediate	Technical
decision or branch coverage	Immediate	Organizational
decode	Immediate	Technical
decrypt	Immediate	Technical
decryption-verification	Immediate	Technical
Decryptor	Immediate	Technical
defect	Immediate	Technical
defect check	Immediate	Technical
defect type	Immediate	Technical
defensive design	Immediate	Technical
degauss	Immediate	Technical
degradation	Immediate	Technical
degraded cybersecurity state	Immediate	Technical
deleted file	Immediate	Technical
denial of service	Immediate	Technical
denial of service (dos)	Immediate	Technical
Denial-of-Service attack	Immediate	Technical
deny by default	Immediate	Technical
Deny list	Immediate	Technical
de-perimeterization	Immediate	Technical
deprecated identifier name	Immediate	Technical
depth	Immediate	Technical
derived piv application	Immediate	Technical

derived piv credential	Immediate	Technical
derived test requirement	Immediate	Technical
descriptive label	Immediate	Technical
design margin	Immediate	Technical
designated approval authority (daa)	Immediate	Organizational
destroy	Immediate	Technical
destroyed state	Immediate	Technical
detailed assessment	Mid	Technical
Detection	Immediate	Technical
deterministic algorithm	Immediate	Technical
developer	Immediate	Technical
DevOps	Short	Sociotechnical
di	Immediate	Technical
diagnostics	Immediate	Technical
dictionary creator	Immediate	Technical
dictionary maintainer	Immediate	Technical
dictionary search	Immediate	Technical
dictionary user	Short	Sociotechnical
differential privacy	Immediate	Technical
diffie hellman (algorithm)	Immediate	Technical
diffie-hellman	Immediate	Technical
digital	Long	Technical
digital asset	Long	Technical
digital identity	Long	Technical
digital rights management	Long	Sociotechnical
digital signature	Long	Technical
digital versatile disc-recordable	Long	Technical
direct black wireline	Immediate	Technical
direct random string	Immediate	Technical
directly identifying variables	Immediate	Technical
discovery	Immediate	Technical
discrete fourier transform test	Immediate	Technical
discrete logarithm cryptography	Immediate	Technical
discussion	Short	Sociotechnical
disinfecting	Immediate	Technical
disinformation	Immediate	Technical
Disinformationists	Immediate	Technical
disintegration	Immediate	Technical
disk image	Immediate	Technical
disposal	Immediate	Technical
disruption	Immediate	Technical
disruptionware	Immediate	Technical
distinguishable information	Immediate	Community/Sector

distributed denial of service	Immediate	Technical
distributed denial of service (ddos)	Immediate	Technical
Distributed Denial-of-Service attack	Immediate	Technical
distributed self-assessment	Immediate	Technical
dlc	Immediate	Technical
dmz	Immediate	Technical
dnp3	Immediate	Technical
dns administrator	Short	Sociotechnical
domain name server	Immediate	Technical
dominance rule	Immediate	Technical
dos	Immediate	Technical
double spend (problem)	Immediate	Technical
downgrading	Immediate	Technical
dpc	Immediate	Technical
drbg mechanism	Immediate	Technical
dual_ec_drbg	Immediate	Technical
duplicate digital evidence	Long	Technical
duty cycle	Immediate	Technical
dvd+r	Immediate	Technical
dvd-rewritable	Immediate	Technical
dvd-rw	Immediate	Technical
dynamic attack surface	Immediate	Technical
ease-of-use	Short	Sociotechnical
e-authentication assurance level	Mid	Technical
education	Mid	Organizational
effective period	Immediate	Technical
electronic credentials	Immediate	Technical
electronic signature	Immediate	Technical
element	Immediate	Technical
element processes	Immediate	Sociotechnical
elliptic curve digital signature algorithm	Long	Technical
emergence	Immediate	Technical
enabling system	Immediate	Technical
encryption certificate	Immediate	Technical
end cryptographic unit (ecu)	Immediate	Technical
end-point protection platform	Immediate	Technical
engineered system	Short	Sociotechnical
enhanced overlay	Immediate	Technical
enhanced security requirements	Immediate	Technical
enrollment	Immediate	Technical
enterprise	Immediate	Technical
enterprise information technology	Mid	Community/Sector
enterprise risk	Mid	Organizational

enterprise service	Immediate	Technical
entropy source	Immediate	Technical
environmental support	Long	Sociotechnical
equivalent process	Mid	Organizational
erm	Immediate	Technical
error	Immediate	Technical
error detection code	Immediate	Technical
examination	Immediate	Technical
examine	Short	Sociotechnical
execute in place	Immediate	Technical
exercise briefing	Short	Sociotechnical
exfiltration	Immediate	Technical
expected output	Immediate	Technical
expert determination	Short	Sociotechnical
exposure	Immediate	Technical
extensible configuration checklist description format (xccdf)	Short	Technical
extension	Immediate	Technical
external coordinator	Immediate	Technical
external security testing	Immediate	Technical
facility	Immediate	Technical
fail safe	Immediate	Technical
fail secure	Immediate	Technical
fail soft	Immediate	Technical
fail to known state	Immediate	Technical
failover	Immediate	Technical
failure	Immediate	Technical
failure access	Short	Sociotechnical
failure control	Immediate	Technical
fair information practice principles	Immediate	Community/Sector
fal	Immediate	Technical
false accept rate (far)	Immediate	Technical
false match rate (fmr)	Immediate	Technical
false negative	Immediate	Technical
false non-match rate	Immediate	Technical
false reject rate (frr)	Immediate	Technical
far	Immediate	Technical
fcb	Immediate	Technical
fckms documentation	Immediate	Sociotechnical
fckms functions	Immediate	Technical
fckms personnel	Immediate	Technical
fckms security domain	Immediate	Technical
fckms services (protections)	Immediate	Technical

fckms service-using organization	Immediate	Technical
feature phone	Immediate	Technical
feature set	Immediate	Technical
features	Immediate	Technical
federal information security	Immediate	Technical
federal information system	Immediate	Organizational
federal information systems security educatorsâ€™ association	Immediate	Technical
federated trust	Immediate	Technical
federation	Immediate	Technical
FedRAMP-compliant	Immediate	Sociotechnical
file name anomaly	Immediate	Technical
file signature anomaly	Immediate	Technical
file slack	Immediate	Technical
file system	Immediate	Technical
fill device	Immediate	Technical
final checklist	Immediate	Technical
fingerprint segmentation	Immediate	Technical
fips pub	Immediate	Technical
firefly	Immediate	Technical
fit for purpose	Immediate	Technical
flaw	Immediate	Technical
focal document	Immediate	Technical
focal document element	Immediate	Technical
focused observation	Immediate	Technical
forced command	Immediate	Sociotechnical
forced ranking optimization	Immediate	Technical
forensic science	Immediate	Technical
formal method	Immediate	Technical
free field	Immediate	Technical
free space	Immediate	Technical
frequency accuracy	Immediate	Technical
frequency drift	Immediate	Technical
fresh	Immediate	Technical
front-channel communication	Immediate	Sociotechnical
frr	Immediate	Technical
full node	Immediate	Technical
functional exercise	Immediate	Technical
fuzz testing	Immediate	Sociotechnical
galois counter mode (algorithm)	Immediate	Technical
Gateway	Immediate	Technical
generation-encryption	Immediate	Technical
geolocation	Immediate	Technical

geometric random variable	Immediate	Technical
georedundancy	Immediate	Technical
global structure/global value	Generational	National/International
global system for mobile communications (gsm)	Immediate	Technical
globally unique identifier	Immediate	Technical
goal	Immediate	Technical
gprs location information	Immediate	Community/Sector
graceful extensibility	Mid	Organizational
graded label	Immediate	Technical
gray market	Long	Community/Sector
graylist	Immediate	Sociotechnical
greatest common divisor	Immediate	Technical
group	Immediate	Technical
group identifier level 1	Immediate	Technical
group order	Immediate	Technical
guest operating system	Immediate	Technical
hackathon	Short	Sociotechnical
hacker	Immediate	Technical
handshake	Immediate	Technical
hardware	Immediate	Technical
hardware asset management	Long	Technical
hardware device	Immediate	Technical
harm	Immediate	Technical
hash chain	Immediate	Technical
hazard	Immediate	Technical
header	Immediate	Technical
health information technology for economic and clinical health act	Long	National/International
health insurance portability and accountability act	Long	National/International
heap	Immediate	Technical
high availability	Immediate	Technical
high impact	Immediate	Technical
high-impact system	Immediate	Technical
high-power transmitter	Immediate	Technical
high-value asset	Immediate	Technical
Honeyport	Immediate	Technical
Honeypot	Immediate	Technical
host operating system	Immediate	Technical
hosted virtualization	Immediate	Technical
hot site	Short	Sociotechnical
https	Immediate	Technical

human user interface capability	Short	Sociotechnical
hypertext transfer protocol secure (https)	Immediate	Technical
hypothesis (null)	Immediate	Technical
ial	Immediate	Technical
ibgp	Immediate	Technical
ict	Immediate	Technical
ict supply chain	Immediate	Community/Sector
ict supply chain risk management	Mid	Organizational
ideal random bitstring	Immediate	Technical
identification and authentication	Immediate	Technical
identified information	Immediate	Community/Sector
identifier cpe name	Immediate	Technical
identifying information	Immediate	Community/Sector
identity and access management	Short	Sociotechnical
identity certificate	Immediate	Technical
identity evidence	Immediate	Technical
identity fraud and identity theft	Immediate	Technical
identity key	Immediate	Technical
identity management system (idms)	Immediate	Technical
identity provider (idp)	Immediate	Technical
identity-based access control	Short	Sociotechnical
identity-based authentication	Immediate	Technical
ied	Immediate	Technical
IIoT	Immediate	Technical
impact level	Immediate	Organizational
impact value	Immediate	Technical
implant	Immediate	Technical
implementation	Short	Technical
incident	Short	Sociotechnical
incident management	Short	Sociotechnical
incineration	Immediate	Technical
incomplete gamma function	Immediate	Technical
independent qualified reviewer	Immediate	Technical
individual accountability	Immediate	Technical
information assurance	Mid	Organizational
information environment	Immediate	Community/Sector
information exchange	Immediate	Community/Sector
information leakage	Immediate	Community/Sector
information life cycle	Immediate	Community/Sector
information relevant to cybersecurity	Immediate	Community/Sector
information resources	Immediate	Community/Sector
information sharing	Mid	Community/Sector

information sharing environment (ise)	Immediate	Community/Sector
information system component	Immediate	Technical
information system life cycle	Immediate	Technical
information system resilience	Immediate	Technical
information systems security (infosec) boundary	Immediate	Technical
information technology	Mid	Community/Sector
information technology product	Mid	Community/Sector
information value	Immediate	Community/Sector
informational label	Immediate	Technical
informative reference developer	Immediate	Technical
informative references	Immediate	Technical
infrastructure as code	Short	Technical
inheritance	Immediate	Technical
initialization vector (iv)	Immediate	Technical
Injury	Immediate	Technical
Injury level	Immediate	Technical
input block	Immediate	Technical
inspection	Immediate	Technical
integrated circuit card id (iccid)	Immediate	Technical
integrated circuit card identification	Immediate	Technical
integrated risk management	Mid	Sociotechnical
integrity	Immediate	Technical
integrity impact	Immediate	Technical
integrity protection	Immediate	Technical
Intellectual property	Immediate	Sociotechnical
intelligence	Immediate	Technical
intelligence activities	Immediate	Technical
interchangeable	Immediate	Technical
interconnection	Immediate	Technical
interconnection security agreement (isa)	Immediate	Technical
interface capabilities	Immediate	Technical
interim approval to operate	Immediate	Technical
interim authorization to test (iatt)	Immediate	Technical
intermittent ad-hoc connection	Immediate	Technical
internal border gateway protocol	Immediate	Organizational
internal control	Immediate	Organizational
international mobile subscriber identity (imsi)	Immediate	National/International
Internet-of-things	Immediate	Technical
invalidate	Immediate	Technical
inventory management	Long	Sociotechnical

IoC	Immediate	Technical
ip	Immediate	Technical
ip security	Immediate	Technical
ipcomp	Immediate	Technical
isa	Immediate	Technical
isao	Immediate	Technical
island of security	Immediate	Technical
isogeny	Immediate	Technical
isolation	Immediate	Technical
issuing source	Immediate	Technical
it asset	Immediate	Technical
ivn	Immediate	Technical
jpeg	Immediate	Technical
judgment value	Immediate	Technical
kas1-basic	Immediate	Technical
kas1-party_v-confirmation	Immediate	Technical
kas2-basic	Immediate	Technical
kas2-party_u-confirmation	Immediate	Technical
kdk	Immediate	Technical
key	Immediate	Technical
key agreement	Immediate	Technical
key bundle	Immediate	Technical
key center	Immediate	Technical
key certification	Long	Technical
key destruction	Immediate	Technical
key escrow system	Immediate	Technical
key establishment	Immediate	Technical
key expansion	Immediate	Technical
key format	Immediate	Technical
key generation	Immediate	Technical
key generation material	Immediate	Technical
key information	Immediate	Technical
key inventory	Immediate	Technical
key life cycle	Immediate	Technical
key list	Immediate	Technical
Key management	Long	Sociotechnical
key management	Long	Sociotechnical
key management components	Long	Technical
key management entity (kme)	Long	Sociotechnical
key recovery	Short	Sociotechnical
key schedule	Immediate	Technical
key share	Immediate	Technical
key signing key (ksk)	Immediate	Technical

key size	Immediate	Technical
key states	Immediate	Technical
key translation center (ktc)	Immediate	Technical
key wrapping	Immediate	Technical
key-center environment	Immediate	Technical
key-derivation key	Immediate	Technical
key-derivation method	Immediate	Technical
key-encryption-key (kek)	Immediate	Technical
key-establishment transaction	Immediate	Technical
keying material	Immediate	Technical
kibi byte	Immediate	Technical
kmi-aware device	Immediate	Technical
kmn	Immediate	Technical
knowledge	Long	Sociotechnical
knowledge levels	Long	Sociotechnical
knowledge-based authentication	Immediate	Technical
koa agent	Immediate	Technical
koa registration manager	Immediate	Technical
l	Immediate	Technical
labeled security protections	Immediate	Technical
laboratory attack	Immediate	Sociotechnical
lan	Immediate	Technical
learning	Immediate	Technical
learning objective	Immediate	Technical
least common multiple	Immediate	Technical
least trust	Immediate	Technical
legacy environment	Long	Technical
level 3	Immediate	Technical
level of assurance	Mid	Technical
life cycle	Immediate	Technical
life cycle security concepts	Immediate	Technical
lightweight directory access protocol (ldap)	Short	Sociotechnical
lightweight node	Immediate	Technical
likelihood	Immediate	Technical
line conditioning	Immediate	Technical
link encryption	Immediate	Technical
linkable information	Immediate	Community/Sector
linked information	Immediate	Community/Sector
local access	Short	Sociotechnical
local element	Immediate	Technical
local registration authority (lra)	Immediate	Organizational
location information (loci)	Immediate	Technical

log archival	Immediate	Technical
log compression	Immediate	Technical
log parsing	Immediate	Technical
log preservation	Immediate	Technical
log reporting	Immediate	Technical
log retention	Immediate	Technical
logic bomb	Immediate	Technical
logical access control system	Immediate	Technical
logical backup	Immediate	Technical
logical perimeter	Immediate	Technical
logical test	Immediate	Technical
logical volume	Immediate	Technical
long runs of ones test	Immediate	Technical
long title	Immediate	Technical
low impact	Immediate	Technical
low probability of detection (lpd)	Immediate	Sociotechnical
low probability of positioning	Immediate	Technical
low-impact system	Immediate	Technical
low-power transmitter	Immediate	Technical
mac key	Immediate	Technical
mac tag	Immediate	Technical
machine-readable	Immediate	Technical
mail transfer agent (mta)	Immediate	Sociotechnical
mail user agent (mua)	Short	Sociotechnical
malicious applet	Immediate	Technical
malicious code	Short	Technical
malicious logic	Immediate	Technical
Malvertising	Immediate	Technical
malware	Immediate	Technical
manage boundaries	Immediate	Technical
manageability	Immediate	Technical
managed environment	Immediate	Technical
management client (mgc)	Long	Sociotechnical
manipulative communications deception	Immediate	Technical
manual cryptosystem	Immediate	Technical
manual remote rekeying	Immediate	Technical
manufacturer usage description (mud)	Immediate	Technical
manufacturing operations	Immediate	Sociotechnical
mapping	Immediate	Technical
margin	Immediate	Technical
market research	Long	Community/Sector
master key	Immediate	Technical

master terminal unit (mtu)	Immediate	Technical
match	Immediate	Technical
matching	Immediate	Technical
materiality	Immediate	Technical
md	Immediate	Technical
means	Immediate	Technical
media access control (mac)	Short	Sociotechnical
medium	Immediate	Technical
medium access control	Short	Sociotechnical
melting	Immediate	Technical
message digest	Immediate	Technical
message inject	Immediate	Technical
microservice	Immediate	Technical
misconfiguration	Immediate	Technical
misnamed files	Immediate	Technical
mission assurance	Mid	Organizational
mission objective	Mid	Organizational
mission resilience	Mid	Organizational
mission-critical element	Immediate	Technical
mission-critical functionality	Immediate	Technical
misuse of controlled unclassified information (cui)	Short	Community/Sector
mms	Immediate	Technical
mobile code risk categories	Short	Technical
model	Immediate	Technical
modern key	Immediate	Technical
moving target defense	Immediate	Technical
mtd	Immediate	Technical
multifactor	Immediate	Technical
multi-level security domain	Immediate	Technical
multiple-center group	Immediate	Technical
multi-releasable	Immediate	Technical
multi-signature	Immediate	Technical
namespace isolation	Immediate	Technical
national comsec incident reporting system (ncirs)	Mid	Organizational
national security information (nsi)	Long	National/International
negative risk	Mid	Organizational
network access	Immediate	Technical
network administrator	Immediate	Technical
network interconnection	Immediate	Technical
network intrusion detection system	Immediate	Technical
network resilience	Short	Organizational

nfiq	Immediate	Technical
nft	Immediate	Technical
nist checklist repository	Immediate	Technical
nist standard	Immediate	Technical
noise injection	Immediate	Technical
no-lone zone (nlz)	Immediate	Technical
non-automated checklist	Immediate	Technical
non-custodial	Immediate	Technical
nonfederal system	Immediate	Technical
nonfungible	Immediate	Technical
nppi	Immediate	Technical
obscured data	Immediate	Technical
observable	Immediate	Sociotechnical
offensive cyberspace operations (oco)	Short	Organizational
official information	Immediate	Community/Sector
offline attack	Immediate	Technical
off-line cryptosystem	Immediate	Technical
olir program	Mid	Technical
on-access scanning	Immediate	Technical
one-time cryptosystem	Immediate	Sociotechnical
one-time pad (otp)	Immediate	Sociotechnical
online attack	Immediate	Technical
open pretty good privacy (openpgp)	Immediate	Technical
open storage	Immediate	Technical
open system	Immediate	Technical
open vulnerability and assessment language (oval)	Short	Technical
OpenIOC	Immediate	Technical
operate & maintain	Immediate	Technical
operation	Immediate	Technical
operational concept	Mid	Organizational
operational environment	Immediate	Organizational
operational margin	Immediate	Organizational
operational phase	Immediate	Organizational
operational resilience	Short	Organizational
operations code (opcode)	Short	Technical
opportunity	Immediate	Technical
opsec	Immediate	Technical
oracle	Immediate	Technical
orphan block	Immediate	Technical
oscillator	Immediate	Technical
other system	Immediate	Technical
otherinput	Immediate	Technical

output space	Immediate	Technical
outside(r) threat	Short	Technical
Overfitting	Immediate	Sociotechnical
overlay network	Immediate	Technical
overt channel	Immediate	Technical
Overwrite	Immediate	Technical
overwrite	Immediate	Technical
package management system	Immediate	Technical
page check	Short	Sociotechnical
pairwise pseudonymous identifier	Immediate	Technical
pairwise trust	Immediate	Technical
paravirtualization	Immediate	Technical
paring code	Short	Technical
passive security testing	Immediate	Technical
passive tag	Immediate	Technical
passwordless	Immediate	Technical
patching	Short	Sociotechnical
ped	Immediate	Technical
Pentester	Immediate	Technical
performance-based	Immediate	Organizational
permanent connection	Immediate	Technical
person	Mid	Sociotechnical
personal accountability	Short	Sociotechnical
personal digital assistant (pda)	Long	Technical
personal firewall	Immediate	Technical
personal information	Immediate	Organizational
personnel security	Mid	Organizational
personnel-security compromise	Short	Organizational
pfs	Immediate	Technical
phase	Immediate	Technical
phishing	Immediate	Technical
physical destruction	Immediate	Technical
physical identifier	Immediate	Technical
pii processing	Immediate	Technical
PIV	Immediate	Technical
piv credential	Immediate	Technical
piv key type	Immediate	Technical
piv visual credential authentication (vis)	Immediate	Technical
pivoting	Immediate	Technical
pkc	Immediate	Technical
plan of action and milestones	Immediate	Technical
plc	Immediate	Technical

Point of presence	Immediate	Technical
portable electronic device (ped)	Immediate	Sociotechnical
portable storage device	Immediate	Technical
positioning	Immediate	Technical
positive control material	Short	Sociotechnical
possession and control of an authenticator	Immediate	Technical
post-market capability	Long	Technical
potential impact	Immediate	Technical
pr	Immediate	Technical
precision	Immediate	Technical
predictability	Immediate	Technical
preimage	Immediate	Technical
preimage resistance	Immediate	Technical
pre-market capability	Long	Technical
prepare for events	Immediate	Technical
preparedness	Immediate	Technical
presentation attack	Immediate	Technical
prf	Immediate	Technical
primary services node (prsn)	Immediate	Technical
prime number	Immediate	Technical
prime number generation seed	Immediate	Technical
primitive	Immediate	Technical
principal authorizing official (pao)	Short	Sociotechnical
privacy impact assessment (pia)	Mid	Sociotechnical
privacy loss	Immediate	Sociotechnical
privacy loss budget	Mid	Technical
private key	Immediate	Technical
private key/private signature key	Immediate	Technical
privilege	Immediate	Sociotechnical
privileged command	Immediate	Technical
probable prime	Immediate	Technical
problem	Immediate	Technical
problematic data action	Immediate	Technical
process outcome	Mid	Organizational
process purpose	Mid	Organizational
processing	Immediate	Technical
product component host	Immediate	Technical
product source node (psn)	Immediate	Technical
profile features	Immediate	Technical
proof of possession (pop)	Immediate	Technical
proof of work consensus model	Immediate	Technical
proper working state	Immediate	Technical

prose checklist	Immediate	Technical
protect & defend	Immediate	Organizational
protected distribution system (pds)	Immediate	Technical
protection	Immediate	Technical
protection bits	Immediate	Technical
provable prime	Immediate	Technical
provider	Immediate	Sociotechnical
Proxyjacking	Immediate	Technical
pseudonymous identifier	Immediate	Technical
pseudorandom	Immediate	Technical
PTaaS	Immediate	Technical
public credentials	Immediate	Sociotechnical
public information	Immediate	Community/Sector
public key	Immediate	Technical
public key cryptographic algorithm	Immediate	Technical
public key infrastructure (pki)	Mid	National/International
public seed	Immediate	Technical
p-value	Immediate	Technical
qemu (quick emulator)	Immediate	Technical
quadratic twist	Immediate	Technical
qualified products list	Immediate	Technical
quality assurance	Mid	Organizational
quality characteristic	Mid	Organizational
quality management	Mid	Organizational
quality property	Mid	Organizational
Quantum computing	Immediate	Technical
questionnaire	Immediate	Technical
random number	Immediate	Technical
random value	Immediate	Technical
randomized hashing	Immediate	Technical
randomized message	Immediate	Technical
randomness source	Immediate	Technical
range	Immediate	Technical
rank (of a matrix)	Immediate	Technical
ransomware	Immediate	Sociotechnical
rapid elasticity	Immediate	Technical
rate	Immediate	Technical
rbac	Mid	Sociotechnical
reader	Immediate	Technical
reader spoofing	Immediate	Technical
real mode	Immediate	Technical
real time reaction	Immediate	Technical
recipient-usage period	Immediate	Technical

reciprocity	Short	Sociotechnical
recommendation	Immediate	Technical
record	Immediate	Technical
records	Immediate	Technical
recovery	Short	Technical
recovery point objective	Short	Technical
recovery time objective	Short	Technical
red key	Immediate	Technical
red team/blue team approach	Mid	Organizational
red wireline	Immediate	Technical
Redaction	Immediate	Technical
redaction	Immediate	Technical
redundancy	Short	Sociotechnical
reference document element	Immediate	Technical
reference monitor	Immediate	Technical
relationship	Immediate	Technical
relationship identifier	Immediate	Technical
relationship style	Immediate	Technical
relationship type	Immediate	Technical
relatively prime	Immediate	Technical
release prefix	Immediate	Technical
relevant event	Immediate	Technical
remote access server	Short	Technical
remote rekeying	Immediate	Technical
replay attack	Immediate	Technical
replay resistance	Immediate	Technical
reporter	Immediate	Technical
representational state transfer (rest)	Immediate	Technical
request for comments	Immediate	Technical
residue	Immediate	Technical
resilience	Short	Organizational
resource	Mid	Organizational
resource pooling	Mid	Organizational
response	Short	Technical
rest	Immediate	Technical
restoration	Immediate	Technical
results	Immediate	Technical
reward system	Immediate	Technical
ripe ncc	Immediate	Technical
risk detail report	Short	Sociotechnical
risk elevation	Mid	Organizational
risk escalation	Mid	Organizational
risk factor	Short	Sociotechnical

risk identification	Short	Organizational
risk management level	Mid	Sociotechnical
risk optimization	Mid	Organizational
risk response	Short	Organizational
risk treatment	Mid	Organizational
risk-based data management	Mid	Organizational
rmf	Immediate	Technical
rng seed	Immediate	Technical
robustness	Immediate	Technical
rogue device	Immediate	Technical
root certificate	Immediate	Technical
root user	Short	Sociotechnical
roots of trust	Immediate	Sociotechnical
rot	Immediate	Technical
route origin attestation	Immediate	Technical
rpo	Immediate	Technical
rs	Immediate	Technical
rtr	Immediate	Technical
rule-based event correlation	Immediate	Technical
run	Immediate	Technical
safety requirements	Immediate	Technical
salt	Immediate	Technical
sap	Short	Organizational
satisfaction	Immediate	Technical
sbom	Mid	Sociotechnical
scada	Immediate	Technical
scap capability	Long	Technical
scap content checklist	Immediate	Technical
scatternet	Immediate	Technical
scenario test	Immediate	Technical
scheduled data transfer	Immediate	Technical
scheme	Immediate	Technical
scheme owner	Short	Sociotechnical
scoping guidance	Immediate	Technical
scrm	Mid	Organizational
sdo	Immediate	Technical
seal of approval	Immediate	Technical
second preimage resistance	Immediate	Technical
secret key (symmetric) cryptographic algorithm	Immediate	Technical
secure communication protocol	Immediate	Sociotechnical
secure communications	Immediate	Technical
Secure erasure	Immediate	Technical

secure state	Immediate	Technical
secure transport	Immediate	Technical
securely provision	Short	Sociotechnical
securely resilient	Short	Sociotechnical
security association database (sad)	Immediate	Technical
security banner	Immediate	Technical
security control and privacy control	Immediate	Technical
security control effectiveness	Immediate	Sociotechnical
security engineering	Immediate	Technical
security information and event management	Short	Sociotechnical
security inspection	Immediate	Technical
security life of data	Immediate	Technical
security marking	Immediate	Technical
security mechanism	Immediate	Technical
security perimeter	Immediate	Technical
security range	Immediate	Technical
security solution	Immediate	Technical
security strength	Immediate	Technical
security-oriented code review	Short	Technical
seed	Immediate	Technical
seed period	Immediate	Technical
selection statement	Immediate	Technical
self-encrypting devices / self-encrypting drives (sed)	Immediate	Technical
sensitive	Short	Sociotechnical
sensitive information	Short	Organizational
sensitivity	Immediate	Technical
serial test	Immediate	Technical
service composition	Immediate	Technical
sha-256	Immediate	Technical
shadow stack	Immediate	Technical
short title	Immediate	Technical
short title assignment requester (star)	Immediate	Technical
short-term stability	Immediate	Technical
shrinkage	Immediate	Technical
side-channel attack	Immediate	Technical
signature	Immediate	Technical
signature generation	Immediate	Technical
signature-in-question	Immediate	Technical
significant consequences	Immediate	Technical
sim	Immediate	Technical
similarity digest	Immediate	Technical

skill	Immediate	Technical
sla	Immediate	Technical
slack space	Immediate	Technical
smart data	Immediate	Technical
smart meter	Immediate	Sociotechnical
sniffer	Immediate	Technical
so	Immediate	Technical
soap header	Immediate	Technical
soap message	Immediate	Technical
sod	Immediate	Technical
solid-state drive	Immediate	Technical
sor	Immediate	Technical
source content	Immediate	Technical
source of randomness	Immediate	Technical
source restriction	Immediate	Technical
source value	Immediate	Technical
sp	Immediate	Technical
space structures	Immediate	Organizational
spam	Immediate	Technical
sparql	Immediate	Technical
special character	Immediate	Technical
specification versioning	Immediate	Technical
sponsor	Immediate	Technical
sponsor (of a certificate)	Immediate	Technical
sponsor (of a key)	Immediate	Technical
spread spectrum	Immediate	Technical
ssh client	Immediate	Technical
ssp	Immediate	Technical
stability	Immediate	Technical
stage	Immediate	Technical
staking	Immediate	Technical
standards developing organization	Long	Organizational
state	Immediate	Technical
strength of function	Immediate	Technical
strength of mechanism (som)	Immediate	Technical
string	Immediate	Technical
strong authentication	Immediate	Technical
subject alternative name	Immediate	Technical
subscriber	Immediate	Technical
suitability and credentialing executive agent	Immediate	Technical
suite a	Immediate	Technical
superior certification authority	Long	Organizational

supplierâ€™s declaration of conformity	Immediate	Technical
supply chain attack	Mid	Organizational
supply chain risk information	Mid	Organizational
support	Short	Sociotechnical
support a security strength	Immediate	Technical
supporting capabilities	Immediate	Technical
supporting parties	Immediate	Technical
supporting services	Immediate	Technical
supportive relationship mapping	Short	Organizational
survivability	Immediate	Technical
suspension	Immediate	Technical
switch	Immediate	Technical
switchport	Immediate	Technical
sybil attack	Immediate	Technical
syncable authenticators	Immediate	Technical
system administration	Immediate	Technical
system developer	Short	Sociotechnical
system flash memory	Immediate	Technical
system high	Immediate	Technical
system integrity	Immediate	Technical
system interconnection	Immediate	Technical
system low	Immediate	Technical
system test	Immediate	Technical
systems development	Mid	Technical
tactical edge	Immediate	Technical
tailored trustworthy space	Immediate	Technical
tamper resistant	Mid	Sociotechnical
target	Immediate	Technical
target data	Immediate	Technical
target value	Immediate	Technical
target vulnerability validation techniques	Short	Sociotechnical
targets	Immediate	Technical
tcp	Immediate	Technical
te	Immediate	Technical
technical reference model (trm)	Mid	Sociotechnical
technical risk	Mid	Organizational
technique	Immediate	Technical
tee	Immediate	Technical
telecommuting	Immediate	Sociotechnical
telework	Immediate	Sociotechnical
test	Immediate	Technical
test action	Immediate	Technical

test guide	Immediate	Technical
test tools	Immediate	Technical
test, training, and exercise (tt&e) plan	Mid	Organizational
three-key triple data encryption algorithm	Immediate	Technical
tier iv checklist	Immediate	Technical
time bomb	Immediate	Technical
timestamp	Immediate	Technical
tls	Immediate	Technical
token	Immediate	Technical
tool configuration	Short	Technical
total risk	Mid	Organizational
TRA	Immediate	Technical
tradecraft identity	Immediate	Technical
trade-off	Immediate	Technical
traffic flow confidentiality (tfc) padding	Immediate	Technical
traffic light protocol	Immediate	Technical
training assessment	Short	Sociotechnical
training effectiveness	Short	Sociotechnical
training key	Short	Sociotechnical
training matrix	Short	Sociotechnical
tranquility	Immediate	Technical
transdisciplinary	Immediate	Technical
transducer capabilities	Immediate	Technical
transforming application	Immediate	Technical
transmission security	Immediate	Technical
transparency	Immediate	Technical
transport mode	Immediate	Technical
Trojanize	Immediate	Technical
trust	Immediate	Sociotechnical
trust anchor	Immediate	Technical
trust list	Immediate	Technical
trust relationship	Immediate	Sociotechnical
trusted	Immediate	Technical
trusted agent (ta)	Immediate	Sociotechnical
trusted association	Immediate	Technical
trusted operating system	Immediate	Technical
trusted timestamp authority (tta)	Immediate	Organizational
trustworthy information system	Immediate	Technical
tsec nomenclature	Immediate	Technical
tsig key	Immediate	Technical
tta	Immediate	Technical
turing complete	Immediate	Technical

two-person control	Immediate	Technical
uniform resource locator	Immediate	Technical
universal description, discovery, and integration (uddi)	Immediate	Sociotechnical
universal integrated circuit card	Immediate	Technical
universal resource identifier	Immediate	Sociotechnical
universal resource locator	Immediate	Technical
update server	Immediate	Technical
upgrade management system	Immediate	Sociotechnical
url	Immediate	Technical
usability	Immediate	Sociotechnical
user activity monitoring	Immediate	Sociotechnical
user principal name	Short	Sociotechnical
valid data element	Immediate	Technical
valid length	Immediate	Technical
validate	Immediate	Technical
validated roa payload	Immediate	Technical
validation	Immediate	Technical
validation model	Immediate	Technical
validator	Immediate	Technical
validity period	Immediate	Technical
value	Immediate	Technical
value string	Immediate	Technical
variable-value configuration	Short	Technical
virtual local area network	Immediate	Technical
virtualized host	Immediate	Technical
voice over internet protocol (voip)	Immediate	Sociotechnical
volatile data	Immediate	Technical
volatile memory	Immediate	Technical
vrp	Immediate	Technical
vulnerability	Short	Technical
vulnerability assessment and management	Short	Sociotechnical
wan	Immediate	Technical
wander	Immediate	Technical
Wargaming	Immediate	Technical
warm site	Short	Organizational
watering hole	Immediate	Technical
watering hole attack	Immediate	Technical
weakest judgment algorithm	Immediate	Technical
weakness	Immediate	Technical
web bug	Short	Technical
web portal	Immediate	Technical

web service interoperability (ws-i) basic profile	Immediate	Technical
well-formed	Immediate	Technical
well-formed cpe name	Immediate	Technical
white team	Short	Sociotechnical
whitelist	Immediate	Technical
whitelisting	Immediate	Technical
wi-fi	Immediate	Technical
wireless application protocol (wap)	Immediate	Technical
wireless markup language	Immediate	Technical
word	Immediate	Technical
workcraft identify	Immediate	Technical
write-blocker	Immediate	Technical
x.509 certificate	Immediate	Technical
xml information security marking (xml-ism)	Immediate	Technical
xquery	Immediate	Technical
zero fill	Immediate	Technical
zero trust	Immediate	Sociotechnical
zero trust architecture	Mid	Organizational
zeroization	Immediate	Technical
zeroize	Immediate	Technical
zone signing key (zsk)	Immediate	Technical
zt	Immediate	Technical