




Article

A Human–AI Collaborative Framework for Cybersecurity Consulting in Capstone Projects for Small Businesses

Ka Ching Chan ^{*,†} , Raj Gururajan [†]  and Fabrizio Carmignani [†] 

School of Business, University of Southern Queensland, Queensland, QLD 4350, Australia;
raj.gururajan@unisq.edu.au (R.G.); fabrizio.carmignani@unisq.edu.au (F.C.)

* Correspondence: kc.chan@unisq.edu.au

† Current address: Springfield Central, Queensland, QLD 4300, Australia.

Abstract: This paper proposes a Human-AI collaborative framework for cybersecurity consulting tailored to the needs of small businesses, designed and implemented within a Master of Cybersecurity capstone program. The framework outlines a structured four-stage development model that integrates students into real-world consulting tasks while aligning with academic and industry objectives. Human–AI collaboration is embedded throughout the process, combining generative AI tools and domain-specific AI agents with human expertise to support the design, delivery, and refinement of consulting resources. The four stages include (1) AI agent development; (2) cybersecurity roadmap creation; (3) resource development; and (4) industry application. Each stage supports both development-oriented outputs—such as templates, training materials, and client deliverables—and research-oriented projects that explore design practices, collaboration models, and consulting strategies. This dual-track structure enables iterative learning and improvement while addressing educational standards and the evolving cybersecurity landscape for small businesses. This framework provides a scalable foundation for capstone-based consulting initiatives that bridge academic learning and industry impact through Human–AI collaboration.



Academic Editors: Martin Gilje Jaatun,
Hanan Hindy and Aunshul Rege

Received: 6 January 2025

Revised: 27 April 2025

Accepted: 30 April 2025

Published: 7 May 2025

Citation: Chan, K.C.; Gururajan, R.; Carmignani, F. A Human–AI Collaborative Framework for Cybersecurity Consulting in Capstone Projects for Small Businesses. *J. Cybersecur. Priv.* **2025**, *5*, 21. <https://doi.org/10.3390/jcp5020021>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: cybersecurity; human–AI collaboration; capstone project

1. Introduction

1.1. Background and Motivation

Small businesses frequently encounter substantial cybersecurity challenges [1] due to their limited access to structured systems and the dedicated resources that larger organisations typically possess. The cybersecurity frameworks intended for larger enterprises often prove too complex and impractical for smaller entities, rendering them susceptible to sophisticated cyber threats [2]. This complexity can lead small businesses to implement ad-hoc security measures, inadvertently creating vulnerabilities in their cyber defences. Moreover, these smaller firms often depend on external service providers or IT support for system management. Yet such reliance can be problematic, as these providers may offer generic monitoring services without delivering tangible, comprehensible data on their effectiveness. Consequently, small business owners may lack the necessary knowledge to evaluate these measures critically, leaving them unable to ascertain the efficacy of their cybersecurity strategies or even to pose relevant inquiries regarding their security posture. These challenges form the core real-world problems that the proposed Human–AI collaborative framework is specifically designed to address.

Human–Artificial Intelligence (Human–AI) collaboration presents a promising solution to address these challenges by developing tailored, accessible tools and frameworks specifically for the needs of small businesses. Additionally, the academic context provides a unique opportunity to engage students in practical, industry-aligned projects through a capstone course, offering them hands-on experience in developing these tools. A structured Research and Development (R&D) agenda can support the development of such resources within a simulated consulting firm, creating sustainable project opportunities for students and aligning educational and industry objectives to achieve real-world impact.

1.2. Purpose and Scope of the R&D Agenda

The goal of this project is to establish a simulated cybersecurity consulting startup capable of delivering comprehensive business and technology consulting services, specifically tailored to the needs of small businesses. To achieve this, the paper proposes an R&D agenda, serving as a structured framework that guides the development of consulting capabilities and resources, with a focus on practical implementation and scalable solutions. Developed within a capstone course in the Master of Cybersecurity program, the agenda provides students with real-world experience while ensuring alignment with the Australian Computer Society (ACS) accreditation standards, as well as the Core Body of Knowledge (CBoK) and Skills Framework for the Information Age (SFIA). The framework is also designed to offer collaborative opportunities for multiple stakeholders—including academics, education designers, program directors, industry partners, the Work-Integrated Learning (WIL) team, and students—to contribute their expertise, provide input, and offer feedback. By integrating educational goals with practical industry applications, the framework's stages and deliverables are structured to support the startup's growth and sustainability, enabling systematic and continuous improvement in consulting practices. Additionally, the agenda is organised to allow students to extend their projects into further research pathways, offering opportunities to pursue Master of Research or PhD studies, contributing to the advancement of cybersecurity knowledge and practices.

This study is guided by three primary research questions (RQs) that address (1) the design of a scalable cybersecurity consulting framework, (2) the integration of Human–AI collaboration to support productivity and scalability, and (3) the alignment of the framework with educational and industry objectives. These questions, along with secondary stage-specific research questions and detailed objectives, are presented in Section 3.

1.3. Structure of the Paper

This paper is organised into seven main sections. Sections 1 and 2 examine foundational methodologies for development-oriented projects and research-oriented projects and Human–AI collaboration, providing context and supporting the selection of approaches for the R&D agenda. Section 3 states the three RQs that shape the initial design and implementation of the R&D agenda and proposes a set of Secondary Research Questions (SRQs) by development stages, and the key objectives of the R&D agenda. Section 4 details the design and development approach of the R&D agenda, discussing core requirements and methodology selections for both development-oriented projects and research-oriented projects, relevant to **RQ1**. Section 5 provides an overview of the four development stages and their deliverables, presenting human-computer collaboration workflows and capstone alignment, directly addressing **RQ2** and **RQ3**. Section 6 reflects on the agenda's impact, scalability, and areas for enhancement, while Section 7 summarises contributions, educational and industry implications, and directions for future research.

2. Literature Review

The R&D agenda for the simulated cybersecurity consulting firm provides a structured framework for developing and evaluating Information Systems (ISs) artefacts tailored to small businesses' needs. To address key challenges in cybersecurity consulting, this agenda relies on a set of foundational methodologies that support both development-oriented and research-oriented projects. When formulating capstone projects, these methodologies should be carefully considered and selected based on their practicality, appropriateness, and ability to meet both industry and academic requirements. Additionally, the feasibility and design of data collection from various stakeholders, such as designers, developers, end-users, and other stakeholders, play critical roles in choosing an appropriate methodology to answer the RQs. Each methodology is thus chosen to facilitate systematic data collection, iterative refinement, and project evaluation, ensuring alignment with the agenda's practical and academic objectives.

2.1. Methodologies for Development-Oriented Projects

Development-oriented projects within the R&D agenda focus on creating practical consulting artefacts, such as tools, templates, and resources, that can be used by small businesses to enhance cybersecurity practices. The primary methodologies for these projects include Design–Build–Test–Learn (DBTL); Agile, Participatory Design (PD); and User-Centered Design (UCD). These methodologies support iterative design and continuous improvement, which are essential for producing effective and user-centric consulting tools.

DBTL is a structured methodology that is widely used in engineering and IS fields, providing a cycle of design, development, testing, and learning that supports rapid prototyping and adaptation [3]. For development-oriented projects, DBTL facilitates quick iterations, allowing teams to refine artefacts based on user feedback and performance testing. This approach aligns well with consulting environments, where solutions must be flexible and responsive to client needs. By iterating through the DBTL cycle, artefacts are continuously refined, enhancing their relevance and applicability for small business cybersecurity.

Agile is another key methodology in the development of consulting tools, emphasising flexibility, incremental improvements, and responsiveness to stakeholder feedback [4,5]. Its iterative nature supports a modular approach to project development, enabling student teams to create artefacts in stages, with each increment building upon previous iterations. This is especially beneficial in short teaching periods, as it allows teams to produce functional artefacts quickly, which can then be further refined based on additional input from clients and stakeholders. Agile's focus on incremental progress ensures that development-oriented projects remain adaptable to evolving cybersecurity needs.

In projects where end-user input is critical, **PD** and **UCD** methodologies are also considered. **PD** directly involves stakeholders in the design process, ensuring that the developed artefacts align with their practical needs and preferences [6,7]. This approach is often viewed as a design approach rather than a strict research methodology [8] and is particularly relevant in consulting contexts, where understanding the unique challenges faced by small business clients can lead to more effective solutions.

UCD, in contrast, encompasses a broad philosophy and a range of methods focused on designing artefacts that prioritise ease of use and accessibility for end-users [9]. UCD describes design processes where end-users influence how the design takes shape; there is a spectrum of ways in which users are involved [10], but the essential principle is that they participate in one way or another. By centering design decisions around user requirements and feedback, UCD enhances artefact usability, making it particularly suitable for tools that small business clients will use independently. Together, PD and UCD help ensure that

artefacts developed within the R&D agenda are not only functional but also accessible and relevant to end-users. A comparison of methodologies for Development-Oriented Projects is provided in Table 1.

Table 1. Comparison of methodologies for development-oriented projects.

| Methodology | DBTL | Agile Development | PD | UCD |
|--------------------------------|---|--|---|---|
| Focus/Goal | Developing and refining practical tools or frameworks through iterative cycles. | Frequent iterations to improve tools with continuous feedback from stakeholders. | Co-creating solutions with stakeholders to ensure relevance and usability. | Designing solutions based on user needs with continuous feedback shaping the outcome. |
| Core Activities | Design, build, test, learn, and refine continuously. | Develop, release, and gather feedback from users iteratively. | Collaborate closely with stakeholders in the design process. | Prototype solutions and refine based on user input. |
| Primary Output | Functional tools, templates, or systems. | Improve tools or processes through iterative releases. | Usable solutions aligned with stakeholder needs. | User-centered artifacts that meet real-world needs. |
| Stakeholder Involvement | Developers only. | Developers and users. | Developers and users. | Developers and users. |
| Role of Reflection | Reflection occurs during the learning phase to drive future improvements. | Reflection occurs at the end of each sprint to adjust the next iteration. | Reflection is continuous, focusing on user involvement to improve outcomes. | Reflection occurs after each user feedback session to refine prototypes. |
| When to Use | When developing and refining tools or frameworks over time. | When quick feedback and iterative improvement are essential. | When stakeholders must be co-creators throughout the process. | When user needs guide the design and development process. |
| Examples of Application | Developing chatbots or cybersecurity tools through iterative prototyping. | Developing tools or templates with frequent sprints and feedback. | Co-creating templates or tools with students and industry partners. | Designing chatbots with continuous client input and feedback. |

2.2. Methodologies for Research-Oriented Projects

Research-oriented projects within the R&D agenda aim to generate insights that contribute to the academic understanding of cybersecurity consulting practices, particularly in the context of Human–AI collaboration. The methodologies supporting these projects include Reflexive Design Science Research (DSR), Action Design Science Research (Action DSR), Action Research (AR), and Mixed-Methods Research. Each methodology provides structured approaches to iterative evaluation, reflection, and refinement of cybersecurity artefacts, tailored to either immediate practical application or deeper academic insights.

Reflexive DSR is a methodology that allows researchers to systematically reflect on the processes and outcomes of artefact design, enhancing understanding through iterative cycles of development and evaluation [11]. Reflexive DSR is particularly suitable for research-oriented projects that require continuous learning from iterative stages, as each cycle contributes to refined insights and improved artefact designs. Unlike reflective practice, which reviews past actions to enhance future outcomes, reflexive practice emphasises ongoing self-awareness of the researcher’s role, assumptions, and influence throughout the design and development process [12]. This self-reflexivity is critical in evaluating cybersecurity artefacts, as it enables researchers to understand how their involvement shapes

outcomes, thus providing deeper insights into the effectiveness of Human–AI collaboration and artefact performance across different contexts.

Action DSR combines elements of AR and DSR, enabling researchers to intervene in real-world settings while systematically evaluating the outcomes of these interventions [11]. Action DSR is particularly effective for research-oriented projects that require hands-on testing of consulting frameworks, allowing students and researchers to observe artefact performance in practice and collect data on their effectiveness [13]. This methodology supports iterative feedback from clients and stakeholders, offering rich insights for continuous improvement and making it well suited to projects where practical engagement with stakeholders is essential [14]. By fostering collaboration between researchers and practitioners, Action DSR enhances the relevance and applicability of research outputs in addressing complex, real-world problems [15].

Action research (AR) is another methodology that supports iterative refinement through active involvement in real-world contexts, often used to improve practices in educational or organisational settings through cycles of planning, action, observation, and reflection [16]. In the context of cybersecurity consulting, AR allows researchers to work collaboratively with clients to identify challenges and co-develop solutions, ensuring that the developed artefacts are relevant to actual business needs. This approach is particularly useful for artefacts tested in live consulting environments, as it enables adaptive design adjustments in response to immediate feedback from stakeholders.

As a widely used research methodology, **Mixed-Methods Research** complements these approaches by integrating qualitative and quantitative data to provide a comprehensive evaluation of artefact performance. For research-oriented projects within the R&D agenda, Mixed-Methods Research enables researchers to capture both measurable outcomes and user perceptions, offering a well-rounded understanding of artefact effectiveness in cybersecurity consulting. While this paper does not report empirical findings, the agenda's structure supports the future collection of both qualitative and quantitative data. For example, qualitative insights may be obtained through student reflection reports, stakeholder interviews, or observations during client simulations. Quantitative data may include metrics such as agent output accuracy, task success rates, or frequency of prompt revisions. These data types are integral to the proposed Mixed-Methods Research approach for evaluating the performance and impact of artefacts developed in Research-Oriented Projects.

In summary, while all three methodologies **Reflexive DSR**, **Action DSR**, and **Action Research** support iterative refinement, they differ in their focus and approach. **Action Research** is primarily reflective and seeks to improve practices through participatory cycles, often directly involving practitioners in the research process [16]. **Action DSR**, in contrast, incorporates design principles to create artefacts that address specific problems while also generating theoretical knowledge, blending practical engagement with structured research objectives [13]. **Reflexive DSR** distinguishes itself by prioritising the researcher's continuous self-reflection on their role in shaping outcomes, thus yielding insights into how researcher bias and assumptions may influence Human–AI interactions and artefact performance [17]. A comparison of methodologies for Research-Oriented Projects is provided in Table 2.

By understanding these distinctions, researchers within the agenda can select methodologies that align with their project goals, whether they seek practical solutions for immediate application or aim to explore the theoretical implications of their design practices. Additionally, the feasibility and design of data collection from various stakeholders, such as designers, developers, end-users, and other stakeholders, play a critical role in choosing an appropriate methodology. The selected approach must support systematic data collection

and align with industry and academic requirements, thereby ensuring that insights from the projects contribute meaningfully to both practical and academic objectives.

Table 2. Comparison of methodologies for research-oriented projects.

| Methodology | Reflexive DSR | Action DSR | Action Research (AR) | Mixed Methods Research (MMR) |
|--------------------------------|--|--|---|---|
| Focus/Core Goal | Reflecting on internal processes to generate generalisable knowledge and insights. | Designing artefacts and engaging stakeholders to refine and apply solutions. | Solving practical problems collaboratively with continuous reflection in action. | Generating comprehensive insights by combining qualitative and quantitative data. |
| Core Activities | Analyse, reflect, and refine internal frameworks and decisions. | Co-create and apply artefacts with stakeholders, reflecting on outcomes. | Apply solutions and reflect within the action cycle. | Collect, integrate, and analyse qualitative and quantitative data. |
| Primary Output | Generalisable frameworks and design principles. | Refined artefacts and practical knowledge. | Practical improvements and process insights. | Balanced insights from both qualitative and quantitative data. |
| Stakeholder Involvement | Developers only. | Developers and users. | Developers and users. | Developers and users. |
| Role of Reflection | Reflection is continuous, focusing on lessons learned from decisions. | Reflection occurs through engagement with stakeholders and application of solutions. | Reflection is integrated into the action process to refine outcomes. | Reflection occurs during data analysis, balancing insights from multiple sources. |
| When to Use | When critical reflection on internal processes is necessary. | When engaging stakeholders to apply and refine solutions. | When solving real-world problems with continuous reflection. | When both qualitative and quantitative data are needed. |
| Examples of Application | Reflecting on the development process of a cybersecurity roadmap. | Testing tools such as chatbots and roadmaps with clients and refining based on feedback. | Implementing and refining frameworks of business strategies for small businesses. | Collecting qualitative and quantitative feedback from clients and students. |

2.3. Human–AI Collaboration

This section explains how Human–AI collaboration enhances cybersecurity consulting practices beyond traditional approaches. Human–AI collaboration is fundamental to the R&D agenda, positioning AI agents as expert partners rather than mere tools. This approach utilises a Mixture of Experts model, where each AI agent is trained in a specific domain, such as NIST frameworks, risk assessment, or regulatory compliance, providing precise, targeted support to human consultants. By carefully selecting the data and information embedded within these AI agents, the R&D agenda ensures that each agent’s knowledge is controlled, reliable, and tailored to deliver accurate insights. For example, a custom Generative Pre-trained Transformer (GPT) Research Integrity Advisor incorporates Australian research guidelines and codes, offering reliable and standard-compliant advice in line with national requirements [18].

Unlike general-purpose Large Language Models (LLMs) like ChatGPT or Gemini, which provide a broad knowledge base, our agenda’s specialised AI agents are developed

using controlled, domain-specific, and company-specific data. This ensures that consulting recommendations remain relevant and consistent. A similar approach is applied in AI-assisted educational design for assessment creation, where collaborative tasks between humans and AI agents span problem-solving, planning, implementation, and refinement stages, enhancing both educational and practical outcomes [19].

The Custom GPT proposed here is a domain-specific AI agent tailored for cybersecurity consulting tasks. It is built on OpenAI's GPT platform, with a system prompt that incorporates detailed cybersecurity frameworks and iteratively refines instructions based on use-case feedback. Although not autonomous, this narrow AI assistant effectively supports consultants by providing consistent, contextualised guidance across consulting stages. It augments human capabilities in information synthesis, document generation, and resource adaptation.

Our agenda emphasises continuous improvement, allowing these agents to be regularly updated and adapted across all project stages from design to evaluation. This adaptability provides small businesses with agile, high-level expertise, particularly in data-intensive tasks like risk analysis. This leaves human consultants free to tackle complex decision-making. In capstone scenarios, students interact with both general LLMs and domain-specific AI agents, gaining critical technical skills through collaborative consulting projects. This environment fosters numerous opportunities for capstone projects that align educational goals with industry-relevant problem-solving, setting the stage for future research.

3. RQs and Objectives

3.1. Primary RQs

This study is guided by three primary RQs that shape the initial design and implementation of an R&D agenda for the simulated cybersecurity consulting firm. These questions focus on establishing a foundation for the agenda that can support rapid, scalable development while addressing educational and industry objectives.

1. **RQ1: What are the core requirements and methodologies for designing an R&D agenda that supports a scalable cybersecurity consulting firm?**
This question examines the essential requirements and methodologies that form the basis of the agenda, ensuring it meets the demands of both consulting and educational contexts.
2. **RQ2: How can the R&D agenda structure human–computer collaboration to enhance productivity and scalability in consulting projects?**
This question investigates how collaborative workflows between human expertise and AI tools, including AI agents, can be embedded within the agenda, aiming to improve productivity and establish a scalable model. It also explores the respective roles of humans and AI agents, examining how each can complement the other to maximise effectiveness in consulting tasks.
3. **RQ3: How can the R&D agenda provide both educational value and industry relevance for capstone projects within a consulting context?**
This question explores how the agenda aligns with capstone objectives and industry requirements, ensuring that students gain practical experience while contributing to the goals of a consulting firm.

These questions address the core aspects of the agenda, creating a structure that supports both present needs and potential future iterations.

3.2. Overview of the Four Development Stages

The R&D agenda for the simulated consulting firm is structured around four core stages, each representing an area for initial exploration and potential future development.

This structure provides a roadmap for establishing consulting capabilities while creating educational opportunities for students. These stages are introduced here to set the context for the SRQs that follow and will be further detailed in the Results section.

1. **Stage 1: AI agent development**—focuses on designing an AI agent trained on cybersecurity frameworks and industry-specific documents to address small business cybersecurity needs, supporting consulting tasks and enhancing productivity.
2. **Stage 2: Cybersecurity roadmap creation**—involves developing a structured cybersecurity roadmap tailored to small businesses, offering a practical guide for these businesses to follow as they advance their cybersecurity practices.
3. **Stage 3: Resource development**—dedicated to creating tools, templates, and training programs that consultants can use and share with clients, establishing an adaptable set of consulting resources.
4. **Stage 4: Industry application**—emphasises the real-world application and testing of resources, gathering feedback from small business clients to inform ongoing refinement and development.

These stages form the preliminary framework of the agenda and serve as the focus of the exploratory SRQs listed below.

3.3. SRQs by Development Stage

To support the initial prototype of the R&D agenda, this study introduces a set of exploratory SRQs. Each SRQ corresponds to one of the agenda's four development stages, focusing on early insights that guide refinement and future studies.

Stage 1: AI Agent Development

- **SRQ1.1: What are the specific design requirements and methodologies for the AI agent to meet consulting and capstone needs?**
- **SRQ1.2: In what ways can the AI agent enhance consulting productivity and accessibility, and what roles does it serve?**
- **SRQ1.3: How will the performance and impact of the AI agents be evaluated in cybersecurity consulting contexts, specifically looking at metrics such as accuracy and user satisfaction, among others, as examples of key indicators?**

Stage 2: Cybersecurity Roadmap Creation

- **SRQ2.1: How does the agenda structure the process for creating customised cybersecurity roadmaps for small businesses?**
- **SRQ2.2: How does the roadmap align with both educational standards and industry requirements?**

Stage 3: Resource Development

- **SRQ3.1: What design principles guide the creation of consulting tools, templates, and training programs within the agenda?**
- **SRQ3.2: How do iterative feedback and human–computer collaboration enhance the usability and effectiveness of consulting resources?**

Stage 4: Industry Application

- **SRQ4.1: What insights are gained from applying the R&D agenda's resources in real-world industry settings, and how effective are they?**
- **SRQ4.2: What lessons can be learned from applying the consulting resources with small business clients?**

These SRQs are intended as exploratory questions to capture initial findings from each stage. Each SRQ has the potential to evolve into a main RQ for future research, allowing for deeper, focused studies that will further validate and refine the R&D agenda.

3.4. Key Objectives of the R&D Agenda

The R&D agenda aims to create a scalable, collaborative framework that meets the immediate needs of a cybersecurity consulting firm while supporting capstone projects in the Master of Cybersecurity program. Its objectives are designed to balance industry relevance and educational value, establishing a foundation for continuous improvement.

1. **Scalable development of consulting resources:** enable the rapid, scalable creation of consulting resources tailored to small business needs, with flexibility for future adjustments.
2. **Facilitating human–computer collaboration:** establish initial workflows that integrate human expertise and AI-driven tools, enhancing productivity and operational efficiency.
3. **Supporting educational and industry alignment:** ensure that the agenda provides educational value aligned with academic standards and real-world industry requirements, offering students hands-on, practical experience.
4. **Continuous improvement and stakeholder collaboration:** foster collaboration among multiple stakeholders, including academics, program directors, industry partners, WIL teams, and students to allow for iterative feedback and refinement.

These objectives support the agenda's initial framework while allowing it to evolve through ongoing stakeholder contributions and future research.

4. Methodology

4.1. Research Design Overview

This section describes the methodologies that guided the initial development of the R&D agenda, providing a sustainable, scalable, multi-year project plan within a capstone course framework. The agenda is structured around four development stages, each encompassing sub-projects designed as either:

1. **Development-oriented projects:** focusing on creating practical consulting tools and resources, and industry applications, or
2. **Research-oriented projects:** generating insights and frameworks for further refinement and academic contributions.

Human–AI collaboration is central to this dual-track approach, enabling participants to leverage AI tools for design, problem-solving, and iterative refinement. The overall methodology aligns with **RQ1** by identifying essential requirements and methodologies, **RQ2** by structuring collaborative workflows, and **RQ3** by aligning the agenda with educational standards in a practical consulting context.

4.2. Methodology Options and Selection for Designing the R&D Agenda

The R&D agenda, an artefact in its own right, was developed to serve as a structured framework for both practical consulting outcomes and academic contributions in the context of a cybersecurity consulting firm. To design and implement this agenda, a combined methodology involving DBTL and Reflexive DSR was selected. This combined approach enables the agenda to achieve both practical outputs and exploratory insights, laying a foundation for iterative improvement. This section elucidates the methodological framework implemented in the R&D agenda, transitioning from the theoretical discussions in the literature review (Sections 2.1 and 2.2) to practical application in a capstone project setting.

The selection of these methodologies was informed by an options analysis that evaluated their alignment with project objectives, stakeholder involvement, and relevance to educational outcomes and industry requirements. While DBTL and Reflexive DSR form the foundational methodologies, the agenda accommodates the integration of additional methodologies such as Agile, Participatory Design (PD), or User-Centered Design (UCD) where specific project requirements or stakeholder feedback indicate their necessity.

This methodological framework not only translates the strategic objectives of the R&D agenda into actionable procedures but also ensures scalability and adaptability. It integrates Human–AI collaboration, aligns with **RQ1** and **RQ2**, and guarantees that the project outcomes are robust, both academically and within industry contexts.

4.3. Human–AI Collaboration for the Cybersecurity Consulting Firm

Human–AI collaboration is integral to the execution of both development-oriented and research-oriented projects. By integrating AI tools with human expertise, the agenda supports collaborative problem-solving, design, and refinement, facilitating productive project execution and continuous improvement. Figure 1 presents a high-level overview of Human–AI collaboration within the simulated consulting firm, illustrating how various stakeholders—including developers, students, and industry partners—contribute to project outcomes.

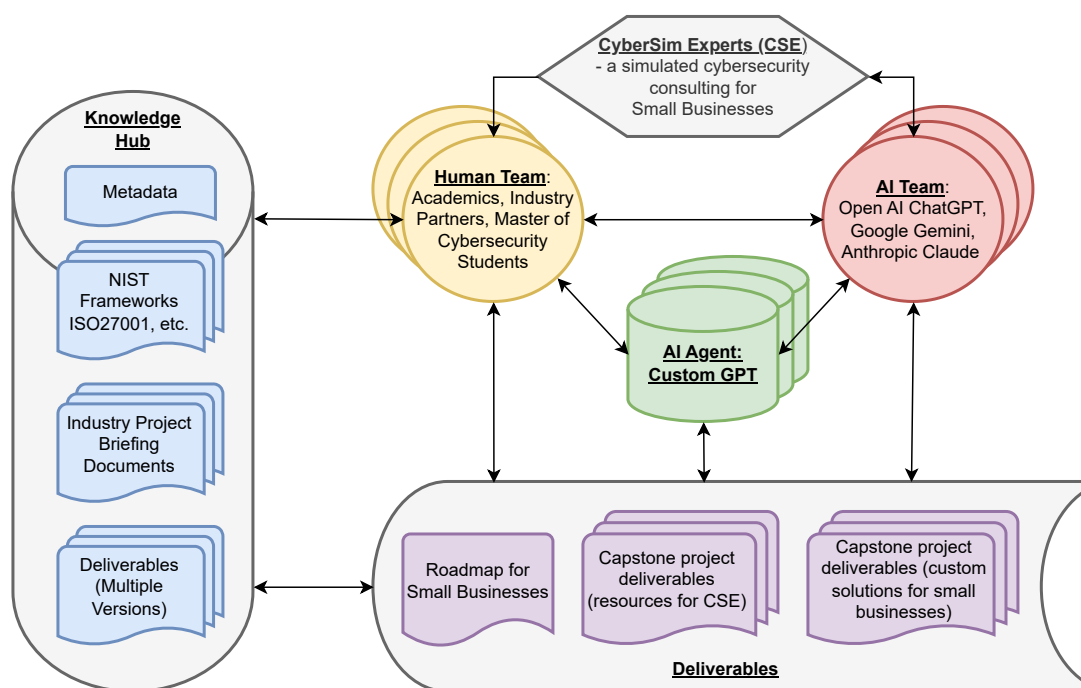


Figure 1. High-level overview of Human–AI collaboration within the simulated cybersecurity consulting firm environment.

This collaborative model enhances both the quality and scalability of deliverables, as AI tools can assist in analysing data, generating insights, and automating tasks. Human–AI collaboration aligns closely with **RQ2**, providing structured workflows that enable iterative feedback and adaptability, which are essential for the dual-track project structure.

4.4. Workflow for Resource Development

The resource development workflow is designed to support iterative creation, feedback, and refinement of consulting resources, with particular emphasis on Human–AI collaboration. Figure 2 details the stages involved in developing resources, from initial

design and prototyping to feedback and finalisation. This workflow incorporates cycles of testing and adaptation to ensure that deliverables remain relevant to small businesses' cybersecurity needs.

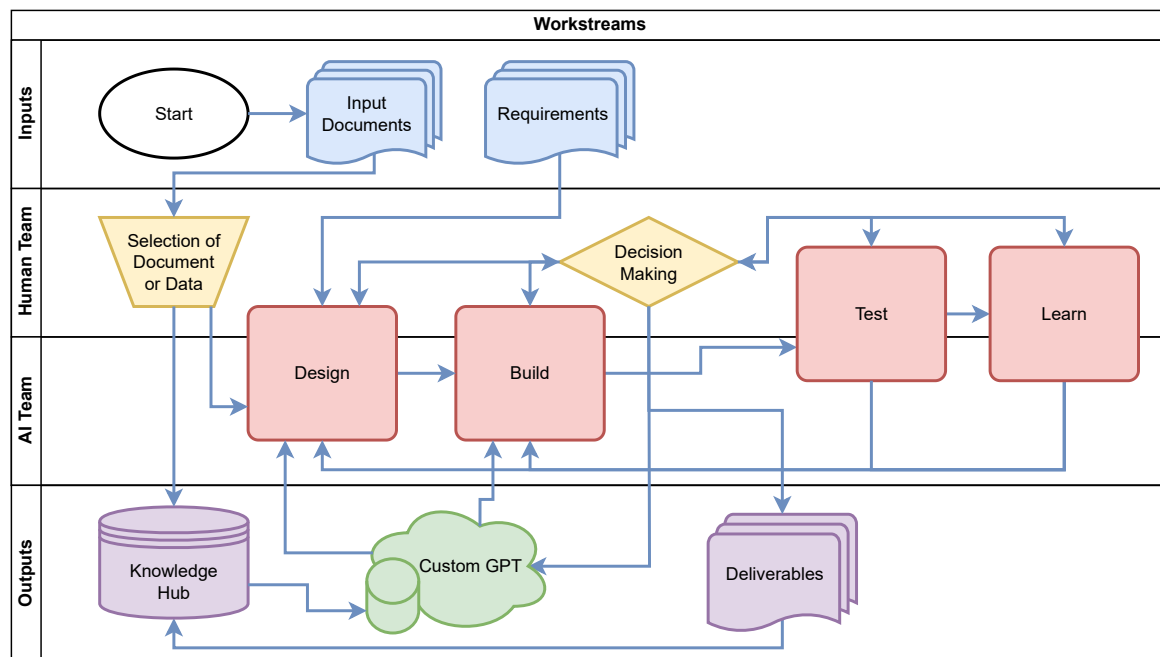


Figure 2. Workflow for Human–AI collaboration in resource development.

By following this structured process, the agenda achieves a balance between rapid prototyping and adaptability, meeting the objectives of both development-oriented and research-oriented projects. This workflow not only supports productivity but also enables ongoing improvements, aligning with the primary RQs by ensuring that resources are responsive to industry and educational requirements.

To ensure that consulting solutions remain current and responsive to the evolving cybersecurity landscape, the framework incorporates a continuous learning mechanism. The prompts used in the Custom GPT agent are updated periodically based on capstone team debriefs, feedback from client engagements, and developments in cybersecurity guidance and threat intelligence. Additionally, consulting resources such as templates and training modules are refined each teaching period, allowing the knowledge base to evolve organically. This adaptive approach ensures that both the human and AI components of the consulting framework remain aligned with industry best practices.

4.5. Summary of Methodology

In summary, the methodology employed in this R&D agenda combines a dual-track structure with selected methodologies, supported by Human–AI collaboration. The default DBTL and Reflexive DSR frameworks enable both practical and research outputs, addressing the primary RQs by ensuring scalable, adaptable, and collaborative project workflows. The dual-track model allows the agenda to deliver consulting resources while providing a foundation for future academic contributions. This integrated approach positions the R&D agenda as a sustainable and flexible model, supporting both capstone learning outcomes and the evolving needs of small business cybersecurity consulting.

5. Design, Structure, and Initial Results of the R&D Agenda

This section presents the design and initial structure of the R&D agenda, outlining each of the four development stages. Each stage is aligned with development-oriented and research-oriented projects, which are structured to address the agenda's primary goals of providing practical consulting resources and generating academic insights. Tables 3 and 4 summarise the core elements of each project type and Figure 3 illustrates the four stages of creating and applying cybersecurity resources for small businesses.

Table 3. Development-oriented projects for the four stages of cybersecurity consulting firm development.

| Aspect | Stage 1: Building the AI Agent | Stage 2: Collaborative Roadmap Creation | Stage 3: Resource Development | Stage 4: Industry Application |
|------------------------|--|---|---|---|
| Purpose | Develop Custom GPT and Knowledge Hub | Develop cybersecurity roadmap and define capstone projects | Create usable resources for the cybersecurity consulting firm | Implement resources in real-world small businesses |
| Human–AI Collaboration | Academics, AI tools, and industry partners collaborate to build the AI Agent | Joint development of the roadmap and capstone projects with AI tools and stakeholders | Joint development of resources (templates, chatbot) with AI tools, students, academics, and industry partners | Real-world application with students working with AI tools and industry partners, applying consulting resources |
| Methodology | DBTL: Iterative development of AI Agent and Knowledge Hub | DBTL: Iterative development of roadmap and capstone projects | DBTL: Iterative development of templates, tools, chatbot with iterative feedback | Agile development: Apply resources with real clients and gather feedback |
| Developers | Academics, Industry Partners | Academics, Students, Industry Partners, University Engagement Team | Academics, Students, Industry Partners | Students, Industry Partners |
| Development Outputs | Custom GPT, Knowledge Hub | Cybersecurity Roadmap, Capstone Project Specifications | Templates, Training Programs, Cybersecurity Advisor Chatbot, Cybersecurity Management Systems | Client-specific cybersecurity documents, processes and systems. Refined resources, updated Knowledge Hub |
| Pedagogy Principles | Constructivism: Learning through collaborative building | Experiential Learning: Hands-on involvement in project creation | Collaborative Learning: Co-creation of knowledge and solutions | Reflective Learning: Gaining insights through reflection and feedback |

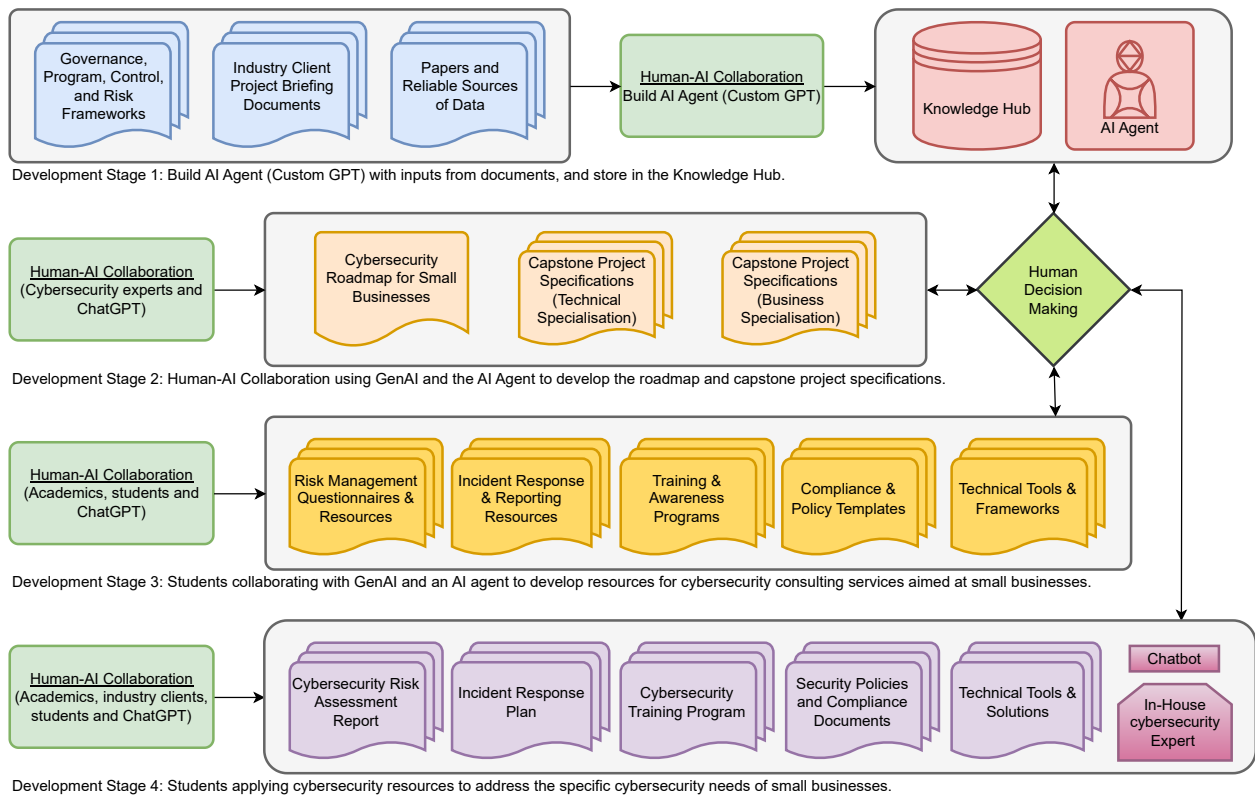


Figure 3. Four development stages of creating and applying cybersecurity resources for small businesses.

Table 4. Research-oriented projects for the four stages of the cybersecurity consulting firm development.

| Aspect | Stage 1: Building the AI Agent | Stage 2: Collaborative Roadmap Creation | Stage 3: Resource Development | Stage 4: Industry Application |
|-------------------------------|--|--|---|---|
| Purpose | Develop frameworks for AI agent creation | Create collaborative frameworks for project development | Generate best practices for collaborative resource development | Produce insights into continuous improvement and feedback loops |
| Example RQs | How can cybersecurity frameworks be integrated into an AI agent for SMEs? | How can students and industry partners collaborate to achieve roadmap development? | What practices ensure effective resource development for SMEs? | How does continuous feedback improve resources and processes over time? |
| Human–AI Collaboration | Reflection on how AI tools and industry data integrate within the AI Agent | Co-creation of collaborative frameworks with input from students and industry | Evaluate collaboration and resource co-creation processes with stakeholders | Continuous feedback for refining resources and improving processes |
| Methodology | Reflexive DSR: Reflection on frameworks and input sources | Action DSR: Co-create roadmap and projects with stakeholder input | Reflexive DSR: Evaluate collaborative process and outcomes | Reflexive DSR: Reflect on feedback and refine resources |
| Researchers | Academics, Industry Partners | Academics, Students, Industry Partners | Academics, Students | Students, Industry Partners, Academics |

Table 4. Cont.

| Aspect | Stage 1: Building the AI Agent | Stage 2: Collaborative Roadmap Creation | Stage 3: Resource Development | Stage 4: Industry Application |
|------------------------------|---|---|---|---|
| Research Outputs (Artefacts) | Framework for AI agent development | Framework for collaborative project creation | Best practices for resource co-creation | Frameworks for continuous feedback and refinement |
| Data Collection | Interviews with academics and industry partners | Workshops and focus groups with students and stakeholders | Observations of collaboration sessions, student reflections | Client feedback, post-implementation reviews, student reflections |

5.1. Overview of the Four Development Stages

The R&D agenda is organised into four sequential stages, each of which builds upon the outputs and insights generated by the previous stage. These stages are designed to support iterative refinement, collaborative engagement, and alignment with small businesses' cybersecurity needs.

1. **AI agent development**—focuses on building a custom AI agent (Custom GPT) that can draw on cybersecurity documents and other inputs to create a knowledge hub. This agent serves as a foundational tool for subsequent stages, providing accessible, relevant information for consulting tasks. Documents such as widely adopted frameworks and best practices [20–26], government resources [27–31] and industry resources [32,33] are categorised and included in the training of the custom AI agent.
2. **Cybersecurity roadmap creation**—utilises Human–AI collaboration, integrating both GenAI and the AI agent to co-develop a cybersecurity roadmap and detailed capstone project specifications. The roadmap is structured as a flexible, five-phase maturity model that can be tailored to a variety of business contexts. Each phase is accompanied by modular templates and diagnostic tools that guide implementation according to the client's size, sector, and technical capacity. This customisation ensures accessibility for low-resourced businesses while preserving scalability for more mature organisations. Additionally, the roadmap is designed to incorporate feedback from industry partners and capstone participants, allowing it to evolve with changing cybersecurity threats and client needs.
3. **Resource development**—involves collaboration between students, GenAI, and the AI agent to create consulting resources aimed at supporting small businesses' cybersecurity needs.
4. **Industry application**—enables students to apply developed resources to real-world cybersecurity challenges faced by small businesses, adapting resources to fit specific needs and gathering feedback for further refinement.

Each stage includes a set of development-oriented and research-oriented projects designed to achieve complementary outcomes. Table 3 provides an overview of development-oriented projects, while Table 4 summarises research-oriented projects within each stage.

5.2. Development-Oriented Projects Across Stages

The development-oriented projects focus on creating practical consulting resources, such as tools, templates, and training programs, that are tailored to meet small businesses' cybersecurity requirements. Table 3 summarises these projects across each stage, detailing aspects such as purpose, methodology, and development outputs.

For each stage, the development-oriented projects aim to achieve a specific purpose:

- **Purpose:** defines the primary goal of each project within the context of consulting needs.
- **Human–AI Collaboration:** highlights how AI tools and human expertise are integrated to enhance productivity and resource creation.
- **Methodology:** specifies the selected methodology (DBTL) for guiding project execution.
- **Developers:** identifies the team members or stakeholders involved in each project.
- **Development outputs:** lists the expected deliverables for each stage, such as consulting tools or templates.
- **Pedagogy principles:** emphasises the educational principles guiding project design to ensure relevance to capstone learning objectives.

5.3. Research-Oriented Projects Across Stages

The research-oriented projects within each stage are designed to generate insights and frameworks that support academic contributions and guide future research. These projects address exploratory RQs and align with the agenda's SRQs. Table 4 provides an overview of research-oriented projects across the four stages, detailing aspects such as purpose, RQs, and data collection methods.

For each stage, the research-oriented projects include the following components:

- **Purpose:** defines the primary objective of each project in terms of generating research insights.
- **Example RQs:** lists specific questions that each project aims to explore, contributing to the agenda's SRQs.
- **Human–AI Collaboration:** describes how Human–AI collaboration supports data collection and analysis processes.
- **Methodology:** specifies Reflexive DSR as the guiding methodology, fostering systematic exploration and reflection.
- **Researchers:** identifies the researchers involved in each stage of the project.
- **Research Outputs:** lists the expected outcomes for each stage, such as frameworks or academic papers.
- **Data collection:** outlines the primary data collection methods, which may include feedback from participants or case study analysis.

5.4. Initial Findings and Practical Insights from Development Stages

The early stages of the R&D agenda have generated initial findings and insights that demonstrate alignment with the primary RQs. Development-Oriented Projects have resulted in practical consulting tools and resources that support small business cybersecurity, while Research-Oriented Projects have produced preliminary insights into effective Human–AI collaboration and iterative refinement.

While Human–AI collaboration offers substantial benefits for productivity and scalability, several limitations must be acknowledged. These include the risk of overreliance on AI-generated content, the opaque reasoning process of LLMs, and the potential for contextual misinterpretation in nuanced cybersecurity scenarios. To mitigate these challenges, we ensure human consultants retain oversight, validate outputs, and iterate on AI-suggested content. Prompt transparency, stakeholder co-design, and feedback loops are embedded in the workflow to ensure the AI's role remains assistive rather than directive.

These initial outcomes provide a foundation for continuous improvement in both consulting practice and academic research, supporting the iterative development of the agenda. They also validate the agenda's structure as a flexible framework that can adapt to the evolving needs of small businesses while contributing to capstone education by engaging students in real-world cybersecurity challenges.

6. Discussion

6.1. Methodology Effectiveness and Alignment (RQ1)

The dual methodology approach, employing DBTL for development-oriented projects and reflexive DSR for research-oriented projects, has proven effective in facilitating rapid prototyping, structured reflection, and continuous feedback integration. The DBTL cycle's adaptability has allowed for iterative improvements across each stage, directly addressing RQ1 by establishing a scalable framework that meets both consulting and capstone goals.

In research-oriented projects, reflexive DSR enabled systematic reflection on exploratory insights, helping researchers evaluate preliminary findings. This structured reflection will be critical in future studies, where it will allow insights to guide subsequent research projects focused on specific components of each development stage. As shown in Tables 1 and 2, the selected methodologies align well with the project goals, yet further studies may incorporate additional methodologies to extend this framework.

6.2. High-Level Impact of Human–AI Collaboration (RQ2)

The integration of Human–AI collaboration across the four development stages demonstrates substantial benefits in productivity, resource creation, and iterative refinement. The AI agent, combined with GenAI tools, augmented human expertise by supporting design, decision-making, and problem-solving. Figure 1 provides a high-level view of this collaboration model, underscoring the roles of AI tools in assisting with resource development and consulting outputs.

However, effective Human–AI collaboration also required careful oversight and regular validation of AI outputs. This oversight emphasises that AI, while valuable for efficiency, still relies on human expertise to ensure that outcomes align with consulting needs. As this agenda progresses, further studies will examine collaboration workflows in greater depth to identify specific AI–human interaction patterns that optimise consulting tasks.

Ethical considerations underpin the design and use of AI-powered tools in this agenda. The Custom GPT agent is positioned explicitly as a support tool rather than a decision-maker, ensuring that students and clients maintain oversight, interpretation, and final responsibility for consulting outcomes. Transparency is embedded through editable outputs, explainable prompts, and clearly defined boundaries for AI contributions. All documents used to train the agent are either publicly available or anonymized with appropriate permissions, ensuring ethical data usage.

Given that these activities are embedded in a university capstone program, all projects operate under the institution's academic integrity and AI usage policies. These policies guide responsible engagement with generative tools and are updated regularly to reflect evolving best practices. Students are also trained as ethical practitioners, critically evaluating AI-generated content, validating outputs, and incorporating stakeholder input throughout the process. This human-in-the-loop approach not only mitigates risks of over-reliance or misuse, but also helps foster trust among small business clients engaging with AI-enhanced solutions.

6.3. Educational and Industry Alignment (RQ3)

The R&D agenda is strategically designed to bridge educational objectives with industry standards and practices, focusing on the cybersecurity needs of small businesses. It structures capstone projects within the Resource Development and Industry Application stages to not only meet industry requirements but also apply relevant cybersecurity frameworks and adhere to best practices. This alignment ensures that students gain hands-on experience that translates directly to real-world consulting scenarios, thereby enhancing their educational experience and providing tangible value to businesses.

An integral part of this alignment involves continuous stakeholder engagement throughout the project lifecycle. Unlike traditional models, in which client input might be limited to the final stages, our framework involves stakeholders from the onset—starting with the initial training of the AI agent using client materials, through to the iterative development and piloting of tools and resources. This approach fosters a co-creative environment where feedback is not only encouraged but is also critical in shaping the project's direction and ensuring the relevance and practicality of the outputs.

The agenda's approach to stakeholder engagement also supports trust-building and outreach beyond the capstone classroom. Clients interact with students in structured and informal settings, providing feedback through interviews, reflections, and satisfaction surveys. This feedback is used to revise consulting resources and retrain the AI agent for future projects. To foster trust, all consulting artefacts are designed with transparency and clarity in mind, using plain-language templates and human-guided workflows. In parallel, the agenda encourages collaboration with external partners—such as government agencies, cyber security associations, and chambers of commerce—to support outreach, improve access to clients, and contribute to the curation of high-quality training materials. These relationships enhance the agenda's long-term sustainability and practical impact.

6.4. Challenges and Lessons Learned

Throughout the study, we encountered challenges such as balancing the dual-track structure of educational goals and industry needs, validating AI outputs, and managing the project focus amid continuous feedback loops. A key lesson learned is the critical importance of early and consistent stakeholder involvement, which has proven essential in adapting to iterative changes and maintaining the relevance of each development stage. Additionally, the necessity of continuously validating AI-generated outputs has underscored the indispensable role of human expertise, particularly in sensitive areas like cybersecurity.

To address these challenges, future work will focus on establishing clearer feedback mechanisms and refining collaborative frameworks among students, developers, and AI tools. These improvements aim to enhance the quality of outputs and ensure they remain aligned with both educational objectives and industry standards. Such measures will reinforce the agenda's capacity for delivering cybersecurity solutions that are both academically robust and practically applicable in real-world settings.

6.5. Future Research Directions and Opportunities for SRQs

This section outlines future improvements and planned expansions of the framework to support long-term evolution and broader application. As an initial, work-in-progress framework, the agenda presents several opportunities for further research aligned with the SRQs. Each development stage—AI agent development, cybersecurity roadmap creation, resource development, and industry application—will be evaluated in greater detail in future studies, allowing each SRQ to evolve into main RQs specific to each stage's focus area.

Future research could investigate the scalability of AI-assisted consulting frameworks, examine the role of Human–AI collaboration in cybersecurity settings, and assess the long-term effectiveness of the resources developed for small businesses. These studies aim to deepen our understanding of how AI capabilities can be leveraged to enhance consulting frameworks and drive meaningful improvements in real-world cybersecurity outcomes.

Designed for scalability and long-term relevance, the R&D agenda's modular structure is ideal for academic adoption and adaptation, fitting various curricula, resource availabilities, and student skill levels. The incorporation of open-source tools, customisable templates, and flexible methodologies facilitates adaptation across a wide range of edu-

cational and consulting scenarios. Additionally, the agenda's Human–AI collaborative model is designed to adapt with emerging technologies and industry practices, supporting ongoing refinement of the framework. This ensures the agenda's sustainability and adaptability over time, making it a robust foundation for continuous innovation in cybersecurity consulting and education.

7. Conclusions

This study presents an initial R&D agenda for a cybersecurity consulting framework that leverages Human–AI collaboration to address the needs of small businesses. Through a dual-track structure of Development-Oriented and Research-Oriented Projects, the agenda aims to guide the design and implementation of practical consulting tools while generating preliminary academic insights, answering RQ1, RQ2, and RQ3. By employing DBTL and Reflexive DSR, this agenda provides a scalable, adaptable model for iterative refinement.

The agenda aligns well with capstone educational objectives by engaging students in real-world applications of cybersecurity, thus preparing them for professional roles in consulting. For consulting firms, this agenda offers a replicable framework that demonstrates how Human–AI collaboration can enhance productivity and resource development. As cybersecurity consulting develops a greater need for efficient, AI-supported solutions, this agenda offers an adaptable, evidence-based approach that can be tailored to various consulting contexts.

As a work in progress, this agenda is positioned for continuous development, with each of its four stages set to be further studied and evaluated in future papers. Upcoming studies will focus on detailed analyses of each stage, transforming SRQs into primary RQs for each focus area. Additionally, future research could investigate the applicability of this framework to larger organisations, explore the evolution of AI capabilities over time, and assess the long-term impact of AI-augmented consulting tools in diverse business environments.

This initial R&D agenda provides a structured approach to integrating Human–AI collaboration within cybersecurity consulting for small businesses. By aligning practical deliverables with academic research across iterative development stages, it offers a replicable model for both educational institutions and consulting firms. The agenda outlines selected methodologies for development and research, contributing to the field of applied Human–AI collaboration and providing valuable insights for capstone courses focused on real-world impact.

This agenda provides a scalable model for delivering cybersecurity consulting solutions that directly address the needs of small businesses—an underserved but high-risk sector. By leveraging Human–AI collaboration in an educational setting, the framework empowers students to develop industry-relevant competencies while creating tangible value for clients. The agenda's dual impact—strengthening small business resilience and enriching cybersecurity education—demonstrates its potential for both immediate application and long-term transformation.

Author Contributions: Conceptualization, K.C.C., R.G. and F.C.; Methodology, K.C.C. and R.G.; Validation, K.C.C. and R.G.; Formal analysis, K.C.C.; Investigation, K.C.C., R.G. and F.C.; Data curation, K.C.C.; Writing—original draft, K.C.C.; Writing—review & editing, K.C.C., R.G. and F.C.; Project administration, F.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data sharing is not applicable.

Acknowledgments: The paper underwent proofreading and refinement using Grammarly and ChatGPT to enhance the language [34].

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Saha, B.; Anwar, Z. A review of cybersecurity challenges in small business: The imperative for a future governance framework. *J. Inf. Secur.* **2024**, *15*, 24–39. [CrossRef]
2. Chidukwani, A.; Zander, S.; Koutsakis, P. A survey on the cyber security of small-to-medium businesses: Challenges, research focus and recommendations. *IEEE Access* **2022**, *10*, 85701–85719. [CrossRef]
3. Elger, D.F.; Beyerlein, S.W.; Budwig, R.S. Using design, build, and test projects to teach engineering. In Proceedings of the 30th Annual Frontiers in Education Conference: Building on A Century of Progress in Engineering Education, Conference Proceedings (IEEE Cat. No. 00CH37135), Kansas City, MO, USA, 18–21 October 2000; Volume 2, p. F3C-9.
4. Chan, F.K.; Thong, J.Y. Acceptance of agile methodologies: A critical review and conceptual framework. *Decis. Support Syst.* **2009**, *46*, 803–814. [CrossRef]
5. Ilieva, S.; Ivanov, P.; Stefanova, E. Analyses of an agile methodology implementation. In Proceedings of the 30th Euromicro Conference, Rennes, France, 3 September 2004; pp. 326–333.
6. Kensing, F.; Blomberg, J. Participatory design: Issues and concerns. *Comput. Support. Coop. Work CSCW* **1998**, *7*, 167–185. [CrossRef]
7. Bossen, C.; Dindler, C.; Iversen, O.S. Evaluation in participatory design: A literature survey. In *Proceedings of the 14th Participatory Design Conference: Full Papers-Volume 1*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 151–160.
8. Spinuzzi, C. The methodology of participatory design. *Tech. Commun.* **2005**, *52*, 163–174.
9. Abras, C.; Maloney-Krichmar, D.; Preece, J. User-centered design. In *Encyclopedia of Human-Computer Interaction*; Bainbridge, W., Ed.; Sage Publications: Thousand Oaks, CA, USA, 2004; pp. 445–456.
10. Vredenburg, K.; Mao, J.Y.; Smith, P.W.; Carey, T. A survey of user-centered design practice. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Minneapolis, MN, USA, 20–25 April 2002; pp. 471–478.
11. Hevner, A.R.; Chatterjee, S. *Design Research in Information Systems: Theory and Practice*; Springer: New York, NY, USA, 2010.
12. Kuechler, B.; Vaishnavi, V. On theory development in design science research: Anatomy of an artifact. *J. Assoc. Inf. Syst.* **2012**, *13*, 490–507.
13. Haj-Bolouri, A.; Purao, S. Action Design Research as a Method-in-Use. *DiVA Portal*. 2018. Available online: <https://www.diva-portal.org/smash/get/diva2:1109694/FULLTEXT01.pdf> (accessed on 5 January 2025).
14. Cronholm, S.; Göbel, H.; Hjalmarsson, A. Empirical evaluation of action design research. In Proceedings of the Australasian Conference on Information Systems, Wollongong, Australia, 5–7 December 2016.
15. Sein, M.K.; Henfridsson, O.; Purao, S.; Rossi, M.; Lindgren, R. Action design research. *MIS Q.* **2011**, *35*, 37–56. [CrossRef]
16. Hendricks, C.C. *Improving Schools Through Action Research: A Reflective Practice Approach*; Pearson: Upper Saddle River, NJ, USA, 2017.
17. Feucht, F.C.; Lunn Brownlee, J.; Schraw, G. Moving beyond reflection: Reflexivity and epistemic cognition in teaching and teacher education. *Educ. Psychol.* **2017**, *52*, 234–241. [CrossRef]
18. Chan, K.C. Building a Research Integrity Advisor with OpenAI GPT Builder. 2024. Available online: <https://ssrn.com/abstract=4884418> (accessed on 5 January 2025).
19. Chan, K.C.; Lokuge, S.; Fahmideh, M.; Lane, M.S. AI-Assisted Educational Design: Academic-GPT Collaboration for Assessment Creation. 2024. Available online: <https://ssrn.com/abstract=4996532> (accessed on 5 January 2025).
20. National Institute of Standards and Technology (NIST). *NIST Cybersecurity Framework (CSF) 2.0*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. Available online: <https://www.nist.gov/cyberframework> (accessed on 5 January 2025).
21. National Institute of Standards and Technology (NIST). *NIST Cybersecurity Framework 2.0: Small Business Quick-Start Guide*; NIST Special Publication 1300; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. [CrossRef]
22. ISO/IEC 27001; Information Security Management Systems—Requirements. International Organization for Standardization: Geneva, Switzerland, 2013. Available online: <https://www.iso.org/standard/54534.html> (accessed on 5 January 2025).
23. Center for Internet Security. *CIS Critical Security Controls v8*; Center for Internet Security: Washington, DC, USA, 2021. Available online: <https://www.cisecurity.org/controls/v8> (accessed on 5 January 2025).
24. ISACA. *COBIT 2019 Framework: Governance and Management Objectives*; ISACA: Rolling Meadows, IL, USA, 2019. Available online: <https://netmarket.oss.aliyuncs.com/df5c71cb-f91a-4bf8-85a6-991e1c2c0a3e.pdf> (accessed on 5 January 2025).
25. Cybersecurity Maturity Model Certification Accreditation Body. *CMC 2.0 Model Overview*; Cybersecurity Maturity Model Certification Accreditation Body. 2023. Available online: <https://www.cyberab.org/> (accessed on 5 January 2025).
26. Australian Cyber Security Centre. *Strategies to Mitigate Cybersecurity Incidents: Essential Eight*; Australian Cyber Security Centre. Available online: <https://www.cyber.gov.au/acsc/view-all-content/essential-eight> (accessed on 5 January 2025).

27. Australian Cyber Security Centre. *Small Business Cyber Security Guide*; Australian Cyber Security Centre. 2023. Available online: <https://www.cyber.gov.au/resources-business-and-government/essential-cybersecurity/small-business-cybersecurity/small-business-cybersecurity-guide> (accessed on 5 January 2025).
28. Australian Cyber Security Centre. *Small Business Cyber Security Checklist*; Australian Cyber Security Centre. 2023. Available online: <https://www.cyber.gov.au/sites/default/files/2023-06/Small%20business%20cyber%20security%20checklist.pdf> (accessed on 5 January 2025).
29. Vazquez, D. *Guide to Cybersecurity for Small and Medium Businesses, Part 1: Understanding Cybersecurity*; USAID: Washington, DC, USA, 2022. Available online: https://pdf.usaid.gov/pdf_docs/PA00ZKCM.pdf (accessed on 5 January 2025).
30. Vazquez, D. *Guide to Cybersecurity for Small and Medium Businesses, Part 2: Protecting Your Business*; USAID: Washington, DC, USA, 2022. Available online: <https://marketlink.org/> (accessed on 5 January 2025).
31. Federal Trade Commission. Cybersecurity for Small Business: Protect Your Small Business. Available online: <https://www.ftc.gov/business-guidance/small-businesses/cybersecurity> (accessed on 5 January 2025).
32. Spanning. *Building Cybersecurity in Small and Midsize Businesses, Cybersecurity Whitepaper*; Spanning. 2018. Available online: <https://www.spanning.com/media/downloads/SB0365-whitepaper-cybersecurity.pdf> (accessed on 5 January 2025).
33. Microsoft. Small Business Resource Center: Discover Tools, Guides, and Expert Advice to Elevate Your Business Success. Available online: <https://www.microsoft.com/en-us/microsoft-365/small-business-resource-center> (accessed on 5 January 2025).
34. Ciaccio, E.J. Use of artificial intelligence in scientific paper writing. *Inform. Med. Unlocked* **2023**, *41*, 101253. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.