PhD in Information Technology and Electrical Engineering
Università degli Studi di Napoli Federico II

PhD Student: Gennaro Esposito Mocerino

Cycle: XXXIX

Training and Research Activities Report

Academic year: 2024-25 - PhD Year: Second

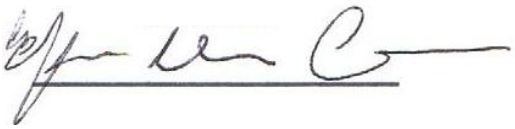

Tutor: prof. Alessio Botta

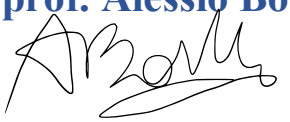

Co-Tutor:

Date: October 31, 2025

Training and Research Activities Report

PhD in Information Technology and Electrical Engineering

Cycle: XXXIX **Author: Gennaro Esposito Mocerino**

1. Information:

- **PhD student: Gennaro Esposito Mocerino** **PhD Cycle:** XXXIX
- **DR number: DR997209**
- **Date of birth: 12/01/1994**
- **Master Science degree: Computer Engineering** **University: Università degli Studi di Napoli Federico II**
- **Scholarship type: no scholarship**
- **Tutor: Prof. Alessio Botta**
- **Co-tutor:**
- **Period abroad:**

2. Study and training activities:

Activity	Type[1]	Hours	Credits	Dates	Organizer	Certificate[2]
How to boost your PhD	Course	18	5	08-15-22-29/01/2025 05-12/02/2025	Prof. Antigone Marino, CNR	Y
Innovation and Entrepreneurship	Course	12	3	5-12-19-26/06/2025 – 23/07/2025	Prof. P. Rippa, Dip. Ing. Industriale	Y
Scientific programming and visualization with Python	Course	30	3	28/01/25 – 11/03/25	Alessio Botta, Stefania Zinno, Giovanni Stanco (organized at DIST)	Y
Perché l'Intelligenza Artificiale crede di fare a meno della teoria linguistica, ma in realtà non potrà farlo	Seminar	2	0.4	20/11/2024	Prof. Franco Cutugno	Y
AI and Enabling Technologies for Social Robots	Seminar	2	0.4	03/12/2024	Prof. Franco Cutugno	Y

Strutture basate su regole e strutture basate su approssimazioni	**Seminar**	**2**	**0.4**	**10/12/2024**	**Prof. Franco Cutugno**	**Y**
URANIA - La Soluzione di E4 per la Cloud-Native AI	**Seminar**	**1**	**0.1**	**16/01/2025**	**E4 Computer Engineering**	**N**
Opportunità e Prospettive dell'AI Generativa nel mondo del Lavoro e della Ricerca	**Seminar**	**3,5**	**0.7**	**29/01/2025**	**ENEA e Università degli Studi di Napoli Federico II**	**Y**
C-CODE OPTIMIZATION for ARM based Embedded System Next Steps	**Seminar**	**4**	**0.8**	**05/03/2025**	**Mario Barbarescchi (Università degli Studi di Napoli Federico II), Salvatore Dello Iacono (Università degli Studi di Brescia), Christian Esposito (Università degli Studi di Salerno) e Giovanni Mazzeo (Università di Napoli Parthenope).**	**N**
5G & Digital transformation a view from an unconventional	**Seminar**	**4**	**0.8**	**14/03/2025**	**5G Academy**	**Y**
How do journals operate	**Seminar**	**1**	**0.2**	**21/05/2025**	**Springer Nature**	**Y**
AI & Scientific Writing	**Seminar**	**1**	**0.2**	**28-05-2025**	**Springer Nature**	**Y**

Benefits of Open Access & SN Author Journey	**Seminar**	**1**	**0.2**	**04/06/2025**	**Springer Nature**	**Y**
Research Integrity	**Seminar**	**1**	**0.2**	**18/06/2025**	**Springer Nature**	**Y**
Real world examples of accepted and rejected submissions	**Seminar**	**1**	**0.2**	**25-06-2025**	**Springer Nature**	**Y**
Trusted Execution Environments for QPUs	**Seminar**	**1**	**0.2**	**27/06/2025**	**prof. Edo Giusto, DIETI-Unina**	**Y**
IEEE Authorship and Open Access Symposium: Tips and Best Practices to Get Published from IEEE Editors	**Seminar**	**1.5**	**0.3**	**15/10/2025**	**IEEE**	**Y**
Guardians or Threats? AI at the Frontlines of Cybersecurity	**Seminar**	**4**	**0.8**	**17/10/2025**	**5G Academy (Giovanni De Bernardo)**	**Y**
AI Powered User interface design	**Seminar**	**4**	**0.8**	**24/10/2025**	**5G Academy (Prof Antonio Origlia)**	**Y**
Quality of services	**Seminar**	**4**	**0.8**	**28/10/2025**	**5G Academy (Prof Federica Battisti)**	**Y**
2025 Spring School on Transferable Skills	**Course**	**9**	**2**	**30-31/10/2025**	**University of Naples Federico II**	**Y**

1) Courses, Seminar, Doctoral School, Research, Tutorship
2) Choose: Y or N

2.1. Study and training activities - credits earned

	Courses	Seminars	Research	Tutorship	Total
Bimonth 1	**0**	**1.2**	**8**	**0.21**	**9.41**
Bimonth 2	**8**	**0.8**	**8**	**0**	**16.8**
Bimonth 3	**0**	**1.6**	**8**	**0.07**	**9.67**
Bimonth 4	**3**	**1.2**	**8**	**0**	**12.2**
Bimonth 5	**0**	**0**	**5**	**0**	**5**
Bimonth 6	**2**	**2.7**	**8**	**0.07**	**12.77**

Total	**13**	**7.5**	**45**	**0.35**	**65.85**
Expected	**30 - 70**	**10 - 30**	**80 - 140**	**0 – 4.8**	

3. Research activity:

During the second year of my PhD, my research focused on expanding and empirically validating the framework introduced in the first year—an Implicit Association Test (IAT) specifically designed to investigate automatic decision processes influencing the susceptibility of the users to phishing attacks [1]. Building upon the platform previously developed, this year's work aimed to conduct a large-scale data collection campaign, refine the experimental design for improved validity, and analyze how implicit associations vary across demographic, occupational, and educational dimensions.

Research Topic and Motivation

Phishing remains one of the most pervasive and costly cybersecurity threats worldwide, responsible for billions in annual losses and countless compromised systems. Yet despite decades of research and extensive awareness programs, people, including trained experts, continue to fall for deceptive messages. Why? Because most defenses and training programs target the wrong part of the mind.

Traditional approaches to phishing prevention [2] are built on slow, deliberate reasoning, the kind of thinking that involves checking the sender's address, hovering over links, or recalling security guidelines. These strategies rely on what Kahneman [3] defines as *System 2*: the rational, analytical mode of thought. However, real phishing attacks rarely allow the luxury of slow reflection. They exploit System 1, our *fast, automatic, and intuitive thinking*, the mental shortcut that drives us to click a link before we even realize why. When an urgent message appears to come from a manager or trusted organization, users act reflexively, driven by psychological cues like urgency, authority, or familiarity. These automatic reactions, not deliberate reasoning, often determine whether a phishing attempt succeeds or fails.

Recent large-scale behavioral research [4] has shown that even trained users can fail to identify phishing attempts when contextual or technical cues are absent. Similarly, predictive models based on behavioral or literacy-related features [5,6] remain constrained to explicit signals and do not capture the automatic processes driving immediate reactions. This gap highlights the need for methodologies capable of uncovering the fast, implicit mechanisms that shape human susceptibility.

That is where this research comes into.

By adapting the Implicit Association Test (IAT), a psychological paradigm introduced by Greenwald et al. [7, 8], this work introduces a novel way to measure how individuals automatically associate phishing-related cues, with trust or suspicion. The IAT is a reaction-time–based task that assesses the strength of automatic mental associations between concepts. In practice, participants are asked to quickly categorize pairs of stimuli, and differences in response times reveal implicit cognitive biases that often operate below awareness. Unlike surveys or behavioral click tests, the IAT captures what people *feel* before they *think*.

Understanding these hidden associations matters because they could reveal who is at risk and why. If we can identify implicit vulnerability patterns, for instance, users who instinctively trust urgent requests or familiar brands, we can design new forms of training and defense mechanisms that target the roots of susceptibility, not just its symptoms.

Ultimately, this research reframes phishing not merely as a technical or educational problem, but as a cognitive one. Measuring implicit associations allows us to move beyond awareness and toward human-

centered cybersecurity, where interventions are tailored to the actual, automatic ways people perceive and react to digital threats.

Methodology

The experimental platform developed during the first year was extended with a parametrized architecture that enabled participants to engage with specific persuasion principles—namely urgency, reciprocity, authority, consistency, liking, and social proof. These principles originate from the seminal work of Robert Cialdini [9], who identified them as the six fundamental mechanisms of social influence frequently exploited in phishing and other social engineering attacks. Each principle represents a distinct psychological lever—urgency pressures users to act quickly, authority invokes compliance toward perceived experts, reciprocity exploits the tendency to return favors, consistency appeals to prior commitments, liking leverages familiarity and positive affect, and social proof relies on the perceived behavior of others.

Participants could select one principle before beginning the test session, allowing the system to dynamically generate a personalized version of the experiment. Based on this selection, the platform automatically constructed trial blocks combining domain name stimuli—either legitimate (e.g., google.com, amazon.com) or phishing-style obfuscated domains (e.g., login-update-paypa1.net, ytube-security.xyz)—with lexical cues representing both the chosen persuasion strategy and its semantic opposite. For instance, in the Urgency condition, words such as immediate, now, and urgent were contrasted with later, calm, and optional.

The experiment followed a seven-block IAT protocol (B1–B7), counterbalanced across participants to avoid order effects. During each block, users rapidly categorized stimuli through binary key presses under conditions of congruency (e.g., phishing + urgency) and incongruency (e.g., phishing + non-urgency). The key assumption is that faster reaction times under congruent conditions reflect stronger implicit associations between phishing cues and the corresponding persuasion principle.

All reaction times and classification errors were recorded with millisecond precision. Data preprocessing adhered to standard IAT analytical conventions, including latency trimming, error penalty correction, and computation of the D-score using the improved algorithm by Greenwald et al. [8]. This score quantifies the strength and direction of implicit associations: higher absolute values indicate stronger automatic cognitive links between phishing-related stimuli and persuasion-based cues.

This modular design allowed for fine-grained analysis of persuasion-specific biases while maintaining flexibility across experimental sessions. The resulting dataset included reaction times, accuracy rates, and demographic information from a diverse participant pool.

Preliminary Results

Preliminary analyses of the data collected so far suggest several emerging patterns. Initial results indicate that IAT-derived measures, such as D-scores, reaction times, and error rates, may hold some predictive value for self-reported phishing exposure. Participants who had previously fallen for phishing attempts tend to display a broader variability in D-scores, suggesting that explicit awareness does not always correspond to implicit resistance.

On average, D-scores were negative (around –0.5), hinting at an implicit tendency to associate phishing-related domains with less persuasive or urgent concepts. This may reflect a latent sensitivity to manipulative cues, although individual differences appear substantial.

Preliminary checks also suggest that task order did not systematically affect response times or accuracy, supporting the internal stability of the IAT design.

Finally, exploratory correlations indicate that demographic and professional factors may moderate implicit associations. For instance, participants with higher education levels generally responded faster and more accurately, whereas those in sectors like Food & Hospitality or Natural Sciences exhibited slightly higher implicit susceptibility. These trends will be further investigated as data collection and analysis progress.

Complementary Work: Analysis of Behavioral Phishing Data through Spamley

In parallel with the IAT-based research, I also continued my work on *Spamley*, a complementary platform developed within my research group to study user behavior in phishing detection tasks. The platform presents participants with a set of ten emails and asks them to decide whether each message is legitimate or a phishing attempt.

During the first year, I updated and optimized the platform to improve its functionality and data collection capabilities, ensuring more reliable and structured response acquisition. This year, I coordinated a new data collection campaign aimed at expanding the participant base and increasing the statistical robustness of the dataset. The newly acquired data, combined with the historical records gathered over previous years, have been systematically processed and analyzed using advanced analytical tools and custom-built scripts.

Preliminary analyses focused on identifying emerging behavioral patterns, error distributions, and temporal decision dynamics across user groups. Particular attention was given to the relationship between demographic variables, confidence ratings, and phishing detection accuracy, which appear to reveal meaningful trends in cognitive and perceptual biases during email evaluation. These preliminary findings are currently being consolidated into a scientific manuscript, which is under preparation for submission to an international conference in the coming months.

In addition, this year we collaborated with the Faculty of Informatics and Computer Science, German International University (GIU), on a complementary study that introduced a predictive framework for user profiling. Using data from *Spamley* [10], the study applied advanced clustering and deep learning techniques to model the relationship between user traits and phishing susceptibility. The resulting framework achieved strong predictive performance, and the outcomes of this collaboration have been published in the *Proceedings of the 11th International Conference on Information Systems Security and Privacy (ICISSP)*, Volume 2.

Implications and Future Work

The overall research carried out during this year reinforces the idea that phishing susceptibility cannot be fully explained through deliberate reasoning or explicit awareness alone. Implicit cognitive mechanisms, automatic associations and heuristic-driven judgments, play a critical role in shaping vulnerability.

By identifying user groups exhibiting higher implicit susceptibility, the findings provide a foundation for developing cognitively informed awareness programs. Such interventions could complement traditional training by addressing automatic decision-making processes, ultimately enhancing cybersecurity resilience at a deeper psychological level.

Future research will expand this framework by integrating implicit measures with behavioral and physiological data to construct more comprehensive user profiles. Additional studies will investigate whether implicit association patterns predict real-world phishing ability, measured through controlled, ethical phishing campaigns. In the long term, this research aims to contribute to a new generation of human-centered cybersecurity solutions that account for the implicit cognitive dimensions of user behavior.

References:

1. *G. Mocerino, C. Velotti, D. Gentile, L. Gallo, A. Botta, and G. Ventre, "Work in progress: Implicit association tests for understanding human factor in phishing Beyond awareness," pp. 519–526, 07 2024*

2. *F. Pietrantonio, A. Botta, G. Ventre, L. Gallo, S. Zinno, L. Mancuso, and R. Presta, "Investigating gaze behavior in phishing email identification," pp. 1–4, 06 2023.*
.
3. *KAHNEMAN, Daniel. Thinking, fast and slow. macmillan, 2011*

4. *Köhler, D., Pünter, W., Meinel, C.: How users investigate phishing emails that lack traditional phishing cues. In: Applied Cryptography and Network Security: 22nd International Conference, ACNS 2024, Abu Dhabi, United Arab Emirates, March 5–8, 2024, Proceedings, Part III. p. 381–411. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-54776-8_15, https://doi.org/10.1007/978-3-031-54776-8_15*

5. *L. Gallo, A. Maiello, A. Botta, and G. Ventre, "2 years in the anti-phishing group of a large company," Computers & Security, vol. 105, p. 102259, 2021.*

6. *L. Gallo, A. Botta, and G. Ventre, "Identifying threats in a large company's inbox," in Proceedings of the 3rd ACM CoNEXT Workshop on Big DAta, Machine Learning and Artificial Intelligence for Data Communication Networks, Big-DAMA '19, (New York, NY, USA), p. 1–7, Association for Computing Machinery, 2019.*

7. *A. G. Greenwald, D. E. McGhee, and J. L. Schwartz, "Measuring individual differences in implicit cognition: the implicit association test.," Journal of personality and social psychology, vol. 74, no. 6, p. 1464, 1998.*

8. *A. Greenwald, B. Nosek, and M. Banaji, "Understanding and using the implicit association test: I. an improved scoring algorithm," Journal of Personality and Social Psychology, vol. 85, pp. 197–216, 08 2003.*

9. *CIALDINI, Robert B.; CIALDINI, Robert B. Influence: The psychology of persuasion. New York: Collins, 2007.*

10. *GALLO, Luigi, et al. The human factor in phishing: Collecting and analyzing user behavior when reading emails. Computers & Security, 2024, 139: 103671.*

4. Research products:

4.1. Gennaro Esposito Mocerino, Claudio Velotti, Danilo Gentile, Luigi Gallo, Alessio Botta, Giorgio Ventre, "Unmasking Implicit Cognitive Biases in Phishing Recognition" in 24th International Conference on Applied Cryptography and Network Security (ACNS), New York City, USA, 2026 – **submitted**

4.2. Wafik P., Botta A., Gallo L., Mocerino G., Herbert C., Annicchiarico I., El Bolock A. and Abdennadher S. (2025). **Enhanced Predictive Clustering of User Profiles: A Model for Classifying Individuals Based on Email Interaction and Behavioral Patterns**. In *Proceedings of the 11th International Conference on Information Systems Security and Privacy - Volume 2: ICISSP*; ISBN 978-989-758-735-1, SciTePress, pages 363-374. DOI: 10.5220/0013302800003899 - **published**

5. Conferences and seminars attended

No Conferences and seminars attended

6. Periods abroad and/or in international research institutions

I have not carried out any activity abroad during my second Ph.D. year.

7. Tutorship

- *Tutoring for Fondamenti di Informatica course, 6h (Bimonth 1)*
- *Tutoring for Reti di Calcolatori course, 2h (Bimonth 3)*
- *Tutoring for Fondamenti di Informatica course, 6h (Bimonth 6)*

8. Plan for year three

The third year of my PhD will be primarily dedicated to data collection and analysis, with the goal of consolidating and extending the preliminary results obtained so far. The main objective is to strengthen the empirical foundation of the research and prepare for the final dissertation.

Core Research Activities

The upcoming year will mainly focus on:

- **Data Collection Campaigns**
  The main effort will involve launching new large-scale data collection campaigns based on the Implicit

Association Test (IAT) framework. These will aim to significantly expand the dataset by including a broader and more diverse sample across demographic, cultural, and professional backgrounds. Data collection will focus on expanding the participant pool across different demographic, cultural, and professional backgrounds, potentially involving new collaborations with external institutions.

- **Data Processing and Analysis**
  A substantial portion of the work will consist of processing and analyzing the growing dataset. Analyses will not only deepen the investigation of patterns already observed but also explore new hypotheses emerging from the expanded data. This may include examining additional behavioral indicators, testing new variables, or applying alternative analytical approaches. The aim is to obtain a comprehensive and validated understanding of implicit cognitive mechanisms related to phishing susceptibility.

Additional and Exploratory Activities

Depending on progress with data collection and analysis, further exploratory activities may include:

- **Preliminary Integration of Multimodal Data**
  The possibility of combining implicit (IAT), behavioral, and demographic data will be explored to assess their joint predictive potential.
- **Initial Machine Learning Experiments**
  Early, small-scale tests may be conducted to evaluate the applicability of machine learning methods for identifying cognitive and behavioral vulnerability patterns.
- **Preparation for Applied, Human-Centric Models**
  Insights gained during this phase will serve as the basis for designing possible human-centered cybersecurity frameworks, to be developed in the final stage of the PhD.

Overall, the third year will focus on expanding data collection, performing comprehensive and diversified analyses, and laying the groundwork for future model development and application.

Collaborations and Projects

The research aims to explore collaborations with institutions in and outside Europe, particularly with universities and research centers specializing in cyberpsychology, human–computer interaction, and cognitive security. These partnerships would enhance the cross-cultural dimension of the work, enabling comparative analyses across different social and technological contexts and fostering the development of a globally relevant framework for human-centric cybersecurity.

Draft Thesis Topic

Tentative Title

"Cognitive and Behavioral Mechanisms in Phishing Susceptibility: An Empirical Analysis"

The thesis will consolidate the theoretical and experimental work conducted through the Implicit Association Test (IAT) framework and complementary behavioral studies (e.g., *Spamley*). The aim is to provide an empirically grounded understanding of phishing vulnerability, integrating explicit and implicit cognitive processes.

Expected Outcomes

By the end of the third year, the research is expected to:

- Complete large-scale data collection campaigns using both the IAT and behavioral experiments, expanding the participant base and strengthening the statistical validity of the results.
- Conduct comprehensive data analyses to validate and refine the existing experimental framework, exploring new variables and behavioral patterns emerging from the enlarged dataset.
- Consolidate preliminary findings into scientific publications, with at least one paper under submission or accepted for presentation at an international conference.
- Lay the groundwork for advanced modeling approaches, such as preliminary machine learning analyses and multimodal data integration, if time and data availability allow.
- Prepare and defend the PhD dissertation, integrating theoretical, experimental, and analytical outcomes into a unified framework.