
UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II
DOTTORATO DI RICERCA / PHD PROGRAM IN
INFORMATION TECHNOLOGY AND ELECTRICAL ENGINEERING

Activities and Publications Report

PhD Student: **Patwardhan**
Narendra

Student DR number: DR996634

PhD Cycle: XXXVIII

PhD Chairman: Prof. Stefano Russo

PhD program student's start date: 01/11/2022

PhD program student's end date: 31/10/2025

Supervisor: Prof. Carlo Sansone

E-mail: carlo.sansone@unina.it

Co-supervisor:

e-mail:

PhD scholarship funding entity:

Student funded by PNRR Program in collaboration with SIMAR GROUP SRL.

General information

Dr. Narendra Patwardhan received in year 2019 the Master of Science degree in Mechanical Engineering from the Michigan Technological University. Within the PhD program in Information Technology and Electrical Engineering, he attended a curriculum in Computer Science. He received a grant from SIMAR GROUP SRL.

Study activities

Attended Courses

Year	Course Title	Type	Credits	Lecturer	Organizer(s)
1st	Statistical Data Analysis for Science and Engineering Research	Ad hoc course	4	Prof. Roberto Pietrantuono	ITEE PhD Program
1st	Introduction to Deep Learning	Ad hoc course	6	Prof. Giovanni Poggi & Prof. Diego Gragnaniello	ITEE PhD Program
1st	Academic Entrepreneurship	Ad hoc course	4	Prof. Pierluigi Rippa	ITEE PhD Program
2nd	Innovation and Entrepreneurship	Ad hoc course	4	Prof. Pierluigi Rippa	ITEE PhD Program

Attended PhD Schools

Year	School title	Location	Credits	Dates	Organizer(s)
1st	Spring School on Transferable Skills	Naples, Italy	2	24/05/2023 - 25/05/2023	Dept. of Pharmacy, UNINA
1st	International Summer School Machine Vision	Padova, Italy	4	04/09/2023 - 08/09/2023	University of Padova

2nd	International Computer Vision Summer School - Computer Vision in the Age of Large Language Models	Catania, Italy	7	07/07/2024 - 13/07/2024	University of Catania
------------	---	----------------	---	-------------------------	-----------------------

Attended Seminars

Year	Seminar Title	Credits	Lecturer	Lecturer affiliation	Organizer(s)
1st	Unleashing the Power of LLMs: a Historical Perspective on Generative AI	0.2	Prof. Carlo Sansone & Prof. Stefano Marrone	UNINA DIETI	ITEE
1st	AI, Robots, and Society: Challenges and Opportunities for Social Innovation	0.2	Prof. Bruno Siciliano	UNINA	ITEE
1st	Symbiotic Control of Wearable Soft Suits for human motion assistance	0.4	Prof. Fanny Ficuciello	UNINA	ITEE
1st	Quantum communications with continuous variables of light	0.3	Prof. Angelica Cacciapuoti	UNINA	ITEE
1st	Building Wearable Assistants with First Person Vision	0.4	ICIAP-23	External	ICIAP-23
1st	Remote Physiological Sensing: State of the Art and Applications	0.4	ICIAP-23	External	ICIAP-23

1st	MATLAB Academic Forum Research and Teaching	0.8	MATLAB	MathWorks	External
2nd	Robotics Meets AI & 5G	0.2	Prof. Bruno Siciliano	UNINA	ITEE
2nd	Hominis	1.0	Prof. Carlo Sansone & Prof. Stefano Marrone	UNINA DIETI	ITEE
2nd	Resource management for mixed-criticality cloud systems	0.2	Prof. Marcello Cinque	UNINA	ITEE
2nd	From ACE Technologies to Sustainable Urban Mobility	0.4	Prof. Stefania Santini	University of Pennsylvania	External
2nd	Learning in nonstationary environments	0.4	Prof. Carlo Sansone	UNINA	ITEE

Research activities

Narendra participated in the research activities of the PICUS Lab at the University of Naples Federico II, focusing on two interconnected research themes: sustainable foundation models and AI-integrated health monitoring systems.

His primary contribution lies in developing the **Hominis** architecture, a human-centric approach to foundation model design that addresses computational sustainability while maintaining performance. He led the pre-training of multiple model variants (8B, 13B, and 15B parameters), developed region-specific models for Italian and Dutch languages, and pioneered model distillation

techniques that reduced parameters from 8B to 0.6B for edge deployment. His work also includes creating comprehensive benchmarking frameworks for evaluating procedural competence and epistemic honesty in language models.

In collaboration with **SIMAR GROUP SRL**, he developed practical AI systems for health monitoring, including containerized voice assistants with natural language interfaces, pseudo-sensor models for non-invasive physiological measurement (heart-rate to blood pressure prediction), and integrated speech-to-text and text-to-speech pipelines for Italian language support. His research bridges theoretical advances in sustainable AI with practical industrial applications, demonstrating expertise in both large-scale distributed training and resource-constrained edge computing.

Tutoring and supplementary teaching activities

No formal tutoring or supplementary teaching activities were performed during the PhD program.

Credits summary

PhD Year	Courses	Seminars	Research	Tutoring / Supplementary Teaching
1 st	20	5.1	35	0
2 nd	11	5.1	45	0
3 rd	0	0	60	0

Research periods in institutions abroad and/or in companies

PhD Year	Institution / Company	Hosting tutor	Period	Activities
2 nd	SIMAR GROUP SRL, Monte Urano (FM), Italy	Industrial supervisor	01/11/2024 - 30/04/2025	Development of AI-integrated health monitoring systems for SIMAR smart chair project. Developed containerized API server for voice assistant, trained pseudo-sensor models for physiological monitoring, and implemented text-to-speech capabilities.
3 rd	RealAI EU, Assen, The	Research collaboratio	01/05/20	Research on sustainable foundation models as part of Hominis project. Trained

	Netherlands	n	25 - 31/10/2025	region-specific language models (Hominis-Italia 8B, Hominis-Dutch 8B), conducted benchmarking studies, and developed educational material generation systems.
--	-------------	---	-----------------	---

PhD Thesis

This PhD thesis addresses critical security, flexibility, and personalization challenges in deploying embodied artificial intelligence in personalized settings. It argues that current agent frameworks are inadequate for safety-critical and privacy-sensitive applications due to architectural deficiencies.

To overcome these limitations, the thesis proposes a vertically integrated architecture founded on three pillars:

- **Transparent Foundation Models:** A family of language models (0.6B, 8B, and 14B parameters) was developed specifically for agentic computation, enabling both high performance and adaptability.
- **Formally Verified Execution Substrate:** A secure, capability-based computational framework built on WebAssembly provides mathematically guaranteed security properties, offering a significant improvement over conventional ad-hoc sandboxing.
- **Gradient-Free Neuroplastic Adaptation:** A novel, on-device learning mechanism inspired by Hebbian learning principles allows agents to continuously adapt to user preferences without computationally intensive training or compromising privacy.

The architecture is validated through its implementation in a “Smart Couch,” a physical prototype integrating health monitoring, comfort optimization, and emergency detection. This real-world application demonstrates the system’s ability to operate with low latency and high accuracy on resource-constrained hardware, while ensuring user privacy through local data processing. The work concludes that a principled, vertically integrated approach can produce trustworthy and adaptable embodied agents, paving the way for more secure and personalized AI systems.

Research products

Research results appear in 3 papers published in international journals, 1 papers published in national journals, and 7 contributions to national conferences.

List of scientific publications

International journal papers

1. N. Patwardhan, S. Marrone, and C. Sansone, Transformers in the Real World: A Survey on NLP Applications, Information, vol. 14(4), 2023.
2. Q. Nasim, N. Patwardhan, T. Maiti, S. Marrone, and T. Singh, VeerNet: Using Deep Neural Networks for Curve Classification and Digitization of Raster Well-Log Images, Journal of Imaging, vol. 9(7), 2023.
3. Marassi, L., Patwardhan, N., Marrone, S., & Sansone, C. Trust in NLP Models: Ethical Considerations on the Reckless Use of AI, Advanced Neural Artificial Intelligence: Theories and Applications, pp. 357-363. Springer Nature Singapore, 2025.

International conference papers

1. Q. Nasim, N. Patwardhan, J. Ali, T. Maiti, S. Marrone, T. Singh, and C. Sansone, Digitizer: A Synthetic Dataset for Well-Log Analysis, 22nd International Conference on Image Analysis and Processing (ICIAP-23), Udine, Italy, September 11-15, 2023.

National conference papers

1. N. Patwardhan, L. Marassi, M. Gravina, A. Galli, M. Zuccarini, T. Maiti, T. Singh, S. Marrone, and C. Sansone, Responsible and Reliable AI at PICUS Lab, Convegno Nazionale CINI sull'Intelligenza Artificiale (Ital-IA 2023), Rome, Italy, May 29-30, 2023.

2. L. Marassi, N. Patwardhan, and F. Gargiulo,
Can Justice Be a Measurable Value for AI? Proposed Evaluation of the
Relationship Between NLP Models and Principles of Justice,
The First Workshop on User Perspectives in Human-Centred Artificial
Intelligence (HCAI4U),
2023.
3. N. Patwardhan, S. Shetye, L. Marassi, M. Zuccarini, T. Maiti, and T. Singh,
Designing Human-Centric Foundation Models,
The First Workshop on User Perspectives in Human-Centred Artificial
Intelligence (HCAI4U),
2023.
4. F. Amato, D. Benfenati, E. Cirillo, G.M. De Filippis, M. Fonisto, A. Galli, S.
Marrone, L. Marassi, V. Moscato, N. Patwardhan, and A. Moccardi,
Advancements and Challenges in Generative AI: Architectures, Applications,
and Ethical Implications,
Convegno Nazionale CINI sull'Intelligenza Artificiale (Ital-IA 2024),
Naples, Italy, May 29-30, 2024.
5. A. Moccardi, E. Cirillo, C. Davino, M. Fonisto, F. Gargiulo, R.C. Ghosh, O.
Gupta, R. Jaiswal, R. La Rovere, L. Marassi, N. Patwardhan, et al.,
Generative AI: Developments, Applications, and Responsible Practices,
Ital-IA 2025,
Trieste, Italy, June 23-24, 2025.
6. A. Galli, F. Amato, D. Carlini, A. Del Prete, S. Dutto, N. Patwardhan, S.
Marrone, V. Moscato, G. Piantadosi, C. Sansone, and M. Valle,
Bridging AI and Industry: Research and Innovation from the CINI-AIIS Lab at
Federico II University,
Ital-IA 2025,
Trieste, Italy, June 23-24, 2025.
7. L. Marassi, N. Patwardhan, S. Marrone, and C. Sansone,
Misinformation and Disinformation in AI: the PICUS Lab Experience,
Ital-IA 2025,
Trieste, Italy, June 23-24, 2025.

Patents and/or spin offs

None.

Awards and Prizes

Best paper award for "Transformers in the Real World: A Survey on NLP Applications" by Information, 2025

Date 25/10/2025

PhD student signature



Supervisor signature


