Carnegie Mellon Portugal

Results of the Carnegie Mellon Portugal Program 2024 Call for PRR Research Grants

Under the 2024 CMU Portugal- PRR Collaborative Research Grants Call, the CMU Portugal Program at CMU will fund twelve (12) new grants, supporting the activities of the participating CMU research teams at Carnegie Mellon University.

A Virtual Reality Game To Transform Attitudes About Climate Change

Principal Investigator at CMU: Laurie M. Heller
Collaborating PRR Project: eGames
Principal Investigator in Portugal: Ana Pires
Supporting researchers: Pedro Campos
Collaborating Institutions: Department of Psychology (CMU); ITI/LARSyS, Instituto Superior Técnico, Universidade de Lisboa; University of Madeira; WOWSYSTEMS – Informática Lda;
Project Abstract: Animal POV: A Climate Challenge is a first-person VR experience that explores how animal embodiment can motivate users to take action to counteract climate change. The PI has developed two versions of a research-oriented game: one allows the user to experience swimming as a dolphin, and the other allows the user to experience flying as a bald eagle. Each experience is composed of three parts that encapsulate how climate change has affected each species. After the player embodies the animal and explores a beautiful and healthy environment, they encounter a climate challenge that disrupts their journey. After navigating through a destroyed environment, the player reaches a restored area. At the end of their journey, they are presented with images of real-world actions, such as habitat restoration efforts, that are combatting the specific climate challenges presented in the game. This educational material is brief and uplifting with the aim of providing the player with motivation to make a difference. Through VR embodiment, the game enables players to see climate change from an animal's perspective, and yet also engage in an autobiographical experience. This transformational game can be used to introduce virtual reality technology and promote sustainable practices. Being inherently fun to play, people may play the game outside of educational settings to expand the reach of its educational impact. We will adapt these games to be appropriate for a population of teenage students and adults in Portugal. This will include translating the instructions, in-game dialogue, and post-game educational materials, and surveys into Portuguese. We will also add references to real-world actions relevant to Portugal. Leveraging existing research being conducted at the ITI-LARSys, we will bring the games to schools of varying socioeconomic status. Our pre-test and post-test surveys will explore how the transformational effect of the game interacts with socioeconomic status. Finally, alternative versions of each game will be created to explore whether a message of despair or hope in a gaming context is more effective at stimulating the desire to learn about and act on climate change.

ACE-FM: Adaptive Conversational Engine powered by Foundation Models

Principal Investigator at CMU: Shinji Watanabe
Collaborating PRR Project: Accelerat.AI
Principal Investigator in Portugal: Alberto Abad
Collaborating Institutions: Language Technologies Institute (CMU); Instituto Superior Técnico (IST)
Project Abstract: This project aims to enhance the capability of speech foundation models to understand long-form, multi-party conversations toward conversational AI. These conversations are influenced by diverse contextual factors, including topics, speakers, environments, speaking styles, and languages. While current speech foundation models show promise in managing this complexity, their large size and lack of flexibility make them ill-suited for adaptively handling such nuanced contexts. Our goal is to develop new foundation models with improved adaptation capabilities, allowing for more efficient and context-aware processing of multi-party conversations. To achieve this, we will first create a benchmark dataset that includes long-form, multi-party conversations and downstream tasks focused on long-form speech recognition and speech summarization. This benchmark will allow for the evaluation of models across key aspects like speaker diarization, topic tracking, and content summarization. We will then develop an adaptive speech foundation model architecture that can generalize prompting abilities to accept a wide range of speech and acoustic contexts, including varied languages and speaking styles. In addition, we will focus on enhancing low-resource language adaptation for target languages and incorporating speech translation and code-switching functionalities. This will enable the model to effectively handle conversations involving multiple languages or dialects, a critical feature for real-world applications. Finally, we are committed to contributing to the research community by making all source code, data, training configurations, and model checkpoints publicly available. This open-science approach will support future work in conversational AI, fostering further innovation in the field.

AI and Human Capital Opportunity Atlas for Portugal

Principal Investigator at CMU: Christophe Combemale,
Collaborating PRR Project: Center for Responsible AI
Principal Investigator in Portugal: Miguel Matos & Paulo Dimas
Supporting researchers: Nikhil George, Ramayya Krishnan, Rahul Telang
Collaborating Institutions: Department of Engineering and Public Policy (CMU); Heinz College of Information Systems and Public Policy (CMU); Católica Lisbon School of Business & Economics; UNBABEL UNIPESSOAL, LDA
Project Abstract: The proposed AI and Human Capital Opportunity Atlas for Portugal would be a data analytics and policy research tool designed to develop a sector-wise AI workforce plan with an initial focus on sectors like Public Administration, Health, Education, Tourism, and Retail. Addressing the critical need for strategic workforce development in an analytics-driven economy, the Atlas would integrate three major data sources to provide actionable insights. Firstly, it would leverage the Community Innovation Survey (CIS), a comprehensive cross-European survey that assesses innovation and the adoption of technologies like artificial intelligence, machine learning, and robotics across various industries. This would allow for benchmarking Portugal's AI integration against other European nations. Secondly, the Atlas would utilize the Quadros de Pessoal microdata to analyze worker mobility patterns within Portugal, identifying how personnel move between firms and sectors, and highlighting effective talent management practices. Lastly, it would incorporate occupation-wise skill information sourced from the US Bureau of Labor Statistics' O*NET \ database, enabling a detailed mapping of the skills and competencies required in different professions. By combining AI adoption data, mobility patterns, and skill requirements, the Atlas aims to identify gaps and opportunities for workforce development, ensuring that training and education initiatives are aligned with the evolving demands of key sectors. This holistic approach would facilitate informed decision-making for

government officials, industry leaders, and educators, fostering collaboration and strategic planning. Ultimately, the proposed AI and Human Capital Opportunity Atlas seeks to drive sustainable innovation, enhance workforce mobility, and strengthen Portugal's competitive edge in the digital economy by providing a robust foundation for talent management and policy development.

AI-mediated Augmented and Assistive Communication Interfaces for ALS: Decoding Movement and Movement Intent Through High-Density EMG and EEG Signals

Principal Investigator at CMU: Pulkit Grover
Collaborating PRR Project: Center for Responsible AI
Principal Investigator in Portugal: Hugo Plácido da Silva
Supporting researchers: Catarina Farinha
Collaborating Institutions: Department Electrical & Computer Engineering (CMU); Instituto de Telecomunicações; Instituto Superior Técnico, Universidade de Lisboa; Unbabel
Project Abstract: The proposed work will optimize muscle and brain machine interfaces (Electromyography – EMG, Electrooculography - EOG, and Electroencephalography - EEG) to advance communication and mobility solutions for individuals with Amyotrophic Lateral Sclerosis (ALS), harnessing the power of AI. Over the course of 1 year, we will conduct a study with 15 healthy volunteers, optimizing Artificial Intelligence (AI) performance on high-density grids. By collecting data on muscle (EMG/EOG) and brain activity (EEG), we aim to establish AI techniques and baseline metrics. This will inform future studies on ALS patients and the creation of next-generation Augmented and Assistive Communication (AAC) interfaces. The volunteers will perform motor tasks, e.g., left/right eye movement, raising eyebrows, or jaw-clenching that typical late-stage ALS patients are able to perform. High-density EMG (on the forehead), EOG, and EEG signals will be recorded. The tasks are designed to a) decode movement using EMG, EOG, and EEG, as well as; b) decode motor intent using EEG (brain activity) alone, prior to any movement. EEG-based decoding movement intent will enable AAC interfaces that remain functional even as muscle activity of ALS patients degenerates. These experiments will allow us to assess the signal quality, reliability, and usability of low-density versus high-density grids in decoding movement and motor intent, providing critical data for refining interface designs. Key milestones include the collection, sharing, and analysis of healthy-volunteer data with our collaborating team in Portugal. At the end of the 1-year period, the resulting optimized system will be transported to Portugal to commence patient studies, and synthesized with Unbabel's HALO software layer, marking the next phase in developing AAC for ALS patients.

CAMCITY – Connected and AI-Driven Mobility for Global Cities

Principal Investigator at CMU: João Barros
Collaborating PRR Project: Route 25: Agenda for Autonomous, Intelligent, Interoperable and Inclusive Mobility
Principal Investigator in Portugal: Ana Aguiar
Supporting researchers: Ricardo Matos
Collaborating Institutions: CMU-Africa; Department of Electrical and Computer Engineering (CMU); Heinz College of Information Systems and Public Policy (CMU); Faculdade de Engenharia da Universidade do Porto ; Nexar
Project Abstract: The CAMCITY project will deliver new AI models and data-driven methodologies to enable connected and automated mobility (CAM) systems to be deployed successfully in diverse global city environments. While existing urban mobility models primarily focus on structured and regulated cities such as Porto, Portugal, and Pittsburgh, USA, African cities like Kigali, Rwanda, present a unique challenge with their mix of formal and informal transport modes, such as bikes, moto-taxis, shared vans, and other vehicles. Leveraging Kigali's rich datasets, including 6 million time-stamped trip logs and 600 hours of vehicle camera

footage, CAMCITY aims to: (a) analyze key differences in mobility patterns and infrastructure between Porto, Pittsburgh, and Kigali, (b) develop AI models to predict CAM system performance in less structured urban environments, and (c) propose solutions that make CAM systems more applicable and inclusive globally. In close collaboration with the Portuguese Route 25 PRR mobilizing agenda, which focuses on autonomous, intelligent, and inclusive mobility, CAMCITY will analyze trac flows, characterize driver behaviors, and predict the performance of CAM technologies across diverse environments. The research outcomes will offer actionable insights to enhance Route 25 CAM technologies, such as vehicle connectivity platforms, automated driving systems and intelligent infrastructure, for operation in more diverse cities thereby contributing to their global relevance and market reach. The project will also foster international collaboration, sending co-supervised researchers between CMU and Portugal to advance the integration of CAM solutions in both developed and developing regions.

Creative Connect: A Data Platform for Climate Action Research

Principal Investigator at CMU: Jessica Hammer
Collaborating PRR Project: eGamesLab
Principal Investigator in Portugal: Valentina Nisi & Nuno J. Nunez
Supporting researchers: Pedro Campos
Collaborating Institutions: Human-Computer Interaction Institute Institute & Entertainment Technology Center (CMU); Interactive Technologies Institute, University of Lisbon; University of Madeira
Project Abstract: Research shows that transformational games can change players' attitudes toward pressing issues such as climate challenges, knowledge about them, and engagement with them. Through this proposal, we are exploring how the impact of the transformations can be translated into out-of-game action. One major challenge in studying the translation process is the lack of a research-integrated climate action platform. Game researchers must choose between high-quality research data (but creating their own climate action context) and authentic climate action (but working with limited data, primarily self-reported). In response to this challenge, we propose to extend our collaboration with the eGames Lab project 1 and connect with Key Conservation, a non-profit that helps conservation organizations gain critical support in real-time from people around the world. Their Creative Connect feature seeks to link consumers of climate content, such as games, with opportunities for action. We will augment this feature to make it research-ready. To ensure our collaboration produces usable knowledge for the field, we will also engage climate partners (e.g. current collaborator Funchal Natural Park in Madeira), and leverage knowledge transfer to partners in the industry (e.g. Wow Systems, Yacooba). We expect several outcomes from this project. First, we will generate new knowledge about how climate researchers conceptualize behavior change. We will learn about potential differences in mental models and user needs by including researchers who already work with games and those who do not. Second, we will provide a new data collection resource for climate researchers working with Unity and a model for extending the Key Conservation platform to support non-Unity games and interventions. Finally, we will transfer what we learn to industry and conservation partners to help other organizations better support data collection for climate change games.

Designing Digital Documentation for Occupational Safety & Health

Principal Investigator at CMU: Jodi Forlizzi, Human-Computer Interaction Institute
Collaborating PRR Project: HfPT - Health from Portugal Prológica Sistemas Informáticos, S.A.
Principal Investigator in Portugal: Ana Correia de Barros
Supporting researchers: Maria Miguel Costa
Collaborating Institutions: Human-Computer Interaction Institute (CMU); Fraunhofer Portugal AICOS; dst center

Project Abstract: This research will examine the challenges of health & safety documentation in the workplace, focusing on the experiences of US-based cleaning & janitorial workers. Despite cleaners' high prevalence of health risks and chronic conditions, traditional reporting mechanisms are difficult to implement, focus on individual versus collective experiences, and hinder intervention unless workers can demonstrate extreme harm. Our project seeks to understand how technology can support the documentation of occupational chronic strain, and how the experience of addressing health in the cleaning sector may inform cross-industry practices. The work will happen in 4 phases: (1) We will first begin with CMU-led ethnographic interviews with cleaning workers in the US. (2) Based on these interviews, CMU and Fraunhofer will work together to develop a range of design probes which reflect workers' experiences with documentation. (3) CMU will then host co-design workshops with US cleaning workers, assisted by a traveling researcher from Fraunhofer (FhP). (4) In our last phase, a researcher from CMU will travel to FhP in Portugal to develop a range of mid-fidelity digital documentation prototypes. In future work, we hope to turn our mid-fidelity prototypes into a series of high-fidelity documentation tools, which we will test with workers on the ground. Taken together, our research sets a strong foundation for the development of new forms of digital documentation, which can inform worker-centered regulatory practices transnationally.

FastCAM: Fast Computation of Optimal Collision-Avoidance Maneuvers

Principal Investigator at CMU: Zachary Manchester,
Collaborating PRR Project: NEURASPACE - AI Fights Space Debris
Principal Investigator in Portugal: Rodrigo Ventura
Collaborating Institutions: Robotics Institute (CMU); Institute for Systems and Robotics – Lisboa, LARSyS, Instituto Superior Técnico
Project Abstract: As Earth orbit becomes increasingly crowded, space-traffic management and collision monitoring are crucial to the continued safe use of space. NEURASPACE is pioneering an AI-based space-traffic monitoring service that can anticipate and identify collision risks between spacecraft and other orbiting objects. This proposal addresses the key next step after a potential collision has been identified: how to maneuver the spacecraft to prevent a collision and ensure safety. This is a complex, multifaceted problem, and we are specifically addressing the following research questions: 1) How do we design maneuvers to account for uncertainty in our knowledge of spacecraft trajectories while still giving strong statistical safety guarantees? 2) How do we design collision-avoidance maneuvers that account for operational constraints and minimize fuel? 3) How do we handle the wide range of spacecraft maneuvering capabilities within a single unified algorithmic framework? 4) How do we perform maneuver computations efficiently and reliably so that they can be automated and performed quickly in time-critical scenarios? To address these questions, we leverage recent advancements in both quasilinear orbital dynamics formulations to efficiently propagate uncertainties, and convex optimization to quickly and reliably compute optimal maneuvers.

Safe Speech Models (SSM)

Principal Investigator at CMU: Bhiksha Raj,
Collaborating PRR Project: Accelerat.AI
Principal Investigator in Portugal: Isabel Trancoso
Supporting researchers: Alberto Abad
Collaborating Institutions:Language Technologies Institute ; Instituto Superior Técnico; INESC-ID Lisboa
Project Abstract: This project proposes to formally study the safety and robustness of speech-enabled generative AI (genAI) models, addressing the risks of harmful outputs such as offensive language, bias, and incitement to harmful behaviours, and investigate solutions for them. With the increasing adoption of these models, including open-source and proprietary versions, a thorough evaluation of their safeguards is critical,

and yet, understudied. Using repositories of harmful behaviours like HarmBench, we will assess these models across three phases: 1) testing their response to vocalized harmful prompts, 2) red-teaming using semantic prompt interventions, and 3) adversarial audio attacks. The project aims to identify vulnerabilities, develop pre-emptive detection mechanisms for harmful responses, and improve model safety through adversarial training and fine-tuning. These findings will guide safer deployment of speech-enabled generative AI models in real-world applications.

SAFE3D-AI: Safety Assurance for 3D Vision AI Systems

Principal Investigator at CMU: David Garlan
Collaborating PRR Project: GreenAuto: Green innovation for the Automotive Industry
Principal Investigator in Portugal: Luís Conde Bento
Supporting reserachers: Abel Mendes
Collaborating Institutions: Software and Societal System Department (CMU); Institute of Systems and Robotics - University of Coimbra; Flowbotic
Project Abstract: The project focuses on applying formal neural network verification to enhance the robustness and safety of the 3D vision system being developed for the Autonomous Mobile Robots (AMRs) in the GreenAuto PRR initiative. The AMR forklift 3D vision system being developed to handle and detect pallet positions and orientations, ensuring safe logistics operations. Given the safety-critical nature of this task, the collaboration between CMU and Portuguese partners aims to integrate formal neural network verification techniques to improve system robustness against real-world anomalies, such as damaged shelves and misaligned pallets, and adversarial attacks. By systematically verifying the neural networks used in the AMR's 3D vision system, the project will ensure that these AI models can safely handle anomalies and meet industrial safety standards. This joint effort between CMU, ISR-UC, and Flowbotic will provide a solution that not only advances scientific knowledge but also has practical applications in logistics automation within the automotive industry.

SSPaM – Sustainable Sensorized Packaging Materials

Principal Investigator at CMU: Carmel S. Majid
Collaborating PRR Project: Embalagem do Futuro (Future of Packaging)
Principal Investigator in Portugal: Mahmoud Tavakoli
Collaborating Institutions: Department of Mechanical Engineering (CMU); Department of Electrical & Computer Engineering (CMU); Instituto de Sistemas e Robótica (ISR), University of Coimbra
Project Abstract: This project will support the mission of the Future of Packaging (Embalagem do Futuro) PRR effort by introducing new classes of thin film circuits and sensors for use as smart packaging materials that are ecologically sustainable. These thin-film materials will be composed of degradable biopolymers and recyclable sensing electronics and function as smart packaging that actively monitors both internal and environmental conditions. Examples of sensing capabilities include temperature, humidity, impact and contact detection, vibration, and changes in orientation/tilt. Combining biodegradability and recyclability with multimodal sensing serves to both reduce waste of the packaging material itself and prevent contents from getting spoiled or damaged during handling and storage. To achieve smart sensing capabilities for future packaging, we will combine gelatin-based biopolymers with biocompatible metal particles (e.g. zinc, molybdenum) to create electronic thin-films that can be used to store food, electronics, and other perishable or delicate products. Depending on the sensing modality, these conductive biopolymer-metal circuits will either function as sensors themselves or be used as flexible (and potentially stretchable) electrical interconnects for surface mounted microelectronic sensing components. This proposed work will build on preliminary collaborative efforts between Prof. Majidi's group at Carnegie Mellon University and Prof. Mahmoud Tavakoli's group at ISR in the University of Coimbra. This includes several joint

publications on printable conductive inks for stretchable electronics and a recent joint study on a fully recyclable smart packaging material using conductive inks made from waterbased polyurethane and silver microflakes. This project will also leverage recent work by the Majidi group on conductive gelatin-based biopolymers, zinc-based biodegradable electronics, and synthesis of conductive inks using metal-organic decomposition.

TOWARD COST-EFFECTIVE PROVABLE SAFETY IN AUTONOMOUS DRIVING

Principal Investigator at CMU: John M. Dolan
Collaborating PRR Project: Project Nr. 29: Route 25
Principal Investigator in Portugal: Susana Sargento
Collaborating Institutions: Robotics Institute (CMU); University of Aveiro
Project Abstract: We propose a new approach for generating depth maps for Autonomous Vehicles (AVs) by integrating the images obtained from deep neural network (DNN) 4D radar detectors with conventional camera RGB images. Our approach introduces a novel pixel positional encoding algorithm inspired by Bartlett's spectrum estimation technique. This algorithm transforms both radar depth maps and RGB images into a unified pixel image subspace, known as the Spatial Spectrum, facilitating the learning of their potential correspondence. Our method effectively leverages high-resolution camera images to train radar depth map generative models, addressing the limitations of conventional radar detectors in complex vehicular environments while sharpening the radar output. We develop spectrum estimation algorithms tailored for radar depth maps and RGB images, a comprehensive training framework for data-driven generative models, and a camera-radar deployment scheme for AV operation. Bypassing expensive and weather susceptible lidars in autonomous vehicle (AV) deployment and for DNN training makes our approach weather-resilient and cost-effective, enabling scalability in AV research and development. Experimental results we obtained so far demonstrate that our approach also outperforms the state-of-the-art (SOTA) by 27.95% in terms of the Unidirectional Chamfer Distance (UCD).