Lifelong Collective Robot Autonomy

Eduardo Sebastián, *University of Cambridge*

I. INTRODUCTION AND MOTIVATION

In multi-robot systems research, in-the-wild scenarios—environmental monitoring, traffic surveillance or emergency construction—serve as quintessential motivating applications. A look at the literature reveals the scale of this focus. For instance, of the 47,382 papers on multi-robot systems in IEEE [12], a 30.14% [13] explicitly mention search and rescue. Yet, no fully autonomous multi-robot system has been deployed in a real-world search and rescue operation. What prevents the adoption of teams of robots in the wild? While many factors contribute, my research focuses on a critical, often overlooked bottleneck: energy autonomy. Multirotors struggle to exceed 30 minutes of flight [10, 52, 8, 45], a window further narrowed by aggressive maneuvers. Ground robots fare slightly better, operating for an hour or two before requiring intervention [2, 49, 3]. These timescales are mismatched with reality: complex human-centric tasks in unstructured environments cannot be resolved in 30 minutes. Therefore, I frame my research around a question: **Can robots help each other to achieve lifelong autonomy?** By lifelong autonomy I mean sustained operation over mission-relevant horizons. I argue that the answer is yes, but only if we embrace **collective intelligence**.

We often treat hardware/software and single/multiple robots as decoupled domains. This leads to algorithmic solutions that are difficult to deploy, ultimately crippling autonomy in robot teams. In contrast, biological collective intelligence achieves lifelong, energy-efficient behavior because computation is distributed across modular, adaptable structures where *form and function* are inextricably linked [11, 15], instead of being confined to a single, monolithic unit [51, 46, 30]. Therefore, my research addresses lifelong autonomy through three pillars: **(i) power system co-design**, optimization and control of the power systems behind the energy supply in robot systems to maximize energy delivery and efficiency; **(ii) automated synthesis of robot designs**, tools that automatically generate robot designs that are modular, decentralized and energy-aware; and **(iii) hardware-aware algorithms**, enabling complex decision-making in teams of robots while remaining computationally lightweight, modular, and responsive to robot constraints.

II. PAST AND CURRENT RESEARCH

Power system co-design. In robotics, the power system that supplies the energy to the robots is often relegated to a secondary role, expected to always provide stable voltage [9]. However, recent research suggests that this is a major barrier to achieving lifelong autonomy [28]. Current trends co-design hardware and controllers [42, 53], integrating electrical and thermal constraints into the control loop via model-based [48] and data-driven battery management [18, 7].

A primary challenge is the need for rapid energy delivery to actuators with minimal energy loss. I have addressed this by developing a nonlinear control strategy for DC-DC converters [35], the electronic component that transfers energy from batteries to the drivers of the actuators. Crucially, the controller is sensorless, adaptive, and implementable on standard microcontrollers, allowing robots to handle fast, changing dynamics without expensive sensing hardware. However, lifelong autonomy also requires a deep understanding of the battery status. I developed a solution for online voltage prediction in energy-constrained installations by tailoring sparse Gaussian Process models [31], enabling fault-tolerant operation in environments where power demand is highly uncertain. More recently, I introduced a method for characterizing power losses through automatic thermal modeling [32]. By identifying optimal discrete-time linear models from temperature profiles, this methodology allows for the accurate monitoring and control of heat dissipation, vital for robots performing aggressive maneuvers or deployed in volatile environmental conditions.

Automated synthesis of robot designs. The design and integration of robot hardware and algorithms is hard because the design process is handcrafted. Recent research in single robot design is looking into methods that formalize the task of generating robot designs that involve computational modules, hardware constraints and task regularities [20, 47, 17, 22]. These span applied category theory [57, 24], evolutionary methods [5, 27, 54], bi-level optimization [55, 50] and simulator-based numerical methods [58, 56], with energy efficiency a primary goal. Yet, current methodologies stop at the individual level, leaving the design of robot collectives underexplored.

My research advances this pillar by providing control- and learning-based frameworks for the automated synthesis of multi-robot behaviors. I introduced Implicit Control [34, 36], which augments the state with the control input to steer heterogeneous, non-cooperative agents whose dynamics are nonlinear in the input—cases where standard controllers fail—treating robots and environment uniformly. Because a synthesized design also fixes its algorithmic modules (Fig. 1b), I learn them directly: physics-informed methods that acquire behaviors from demonstrations [37] or reward signals [41] and scale to large collectives. These methods impose port-Hamiltonian structures on the policy representation that exploits physical constraints such as energy conservation and kinematic limits. The synthesized designs demonstrate zero-shot sim-to-real transfer and robustness in communication.

Hardware-aware algorithms. With an efficient energy supply, and a multi-robot design that integrates hardware constraints, the next step is to devise deployable algorithms. Emerging methodologies are focusing on event-triggered com-

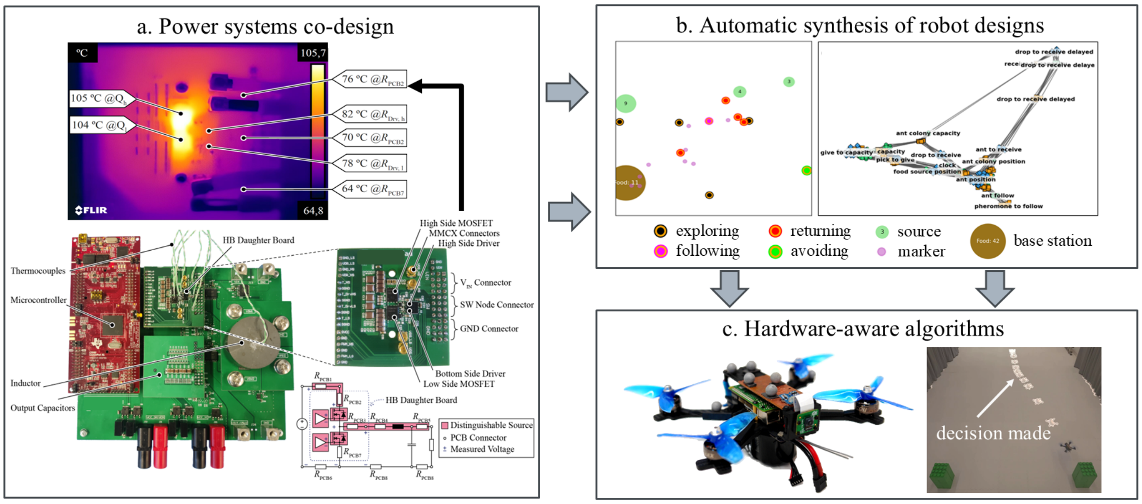

Fig. 1: **My long-term goal, a unified view of lifelong robot autonomy through the lens of collective intelligence**. **a)** Power systems co-design integrates thermal and electrical optimization, estimation and control to maximize on-board energy efficiency. **b)** Given efficient hardware modules, we leverage principles from collective intelligence to synthesize modular robot designs as graphs from high-level specifications (”solve multi-robot foraging”) that integrate environment (food sources), robot embodiments (robots and stigmergy markers) and algorithmic computation (exploration, food return, stigmergy marker following, and collision avoidance). **c)** The realization of the computational modules is achieved by hardware-aware algorithms that enable low-cost decision-making.

munications [16, 6] and communication-free protocols that only rely on onboard sensors [4, 23] (a reliable communication infrastructure is expensive in terms of energy), distilled neural-networks and highly optimized controllers to enable low-cost decision-making [19, 26], and bio-inspired architectures that minimize energy consumption and latency [1, 43].

In this third pillar, I introduced an accelerated ADMM-based gradient tracking method for distributed optimization [40] that achieves faster convergence rates than standard first-order methods with the same communication and computation cost. Furthermore, I designed the first event-triggered certifiable optimal distributed Kalman filter [38], ensuring optimal state estimation while significantly reducing the number of messages exchanged between robots. This is combined with a novel discrete-time consensus protocol [39] that enhances convergence speed through a multi-stage structure that fully exploits the layout of communication messages. Furthermore, I developed an adaptive optimal collision avoidance method driven by an opinion dynamics model [21], in which robots can adapt their safety margins based on the behavior of nearby sensed agents, ensuring safe operation in dense scenarios without communications and resourceful computation. More recently, I have developed the first deployable model for prompting robot teams through natural language instructions in low-cost platforms [29]. The approach distills the reasoning capabilities of LLMs in a tiny recurrent model that modulates goal-oriented policies. The result is a decentralized, modular system that allows humans to convey semantic information to the robots, that reason and collaborate in real-time.

III. FUTURE WORK

Having established fundamental methodologies for hardware-aware estimation, design and control of multi-robot systems, I have started to show how an energy perspective can be integrated into robot collectives. I am now moving to a unified view of robot design that treats collective intelligence as the primary mechanism for lifelong autonomy. Integrated, the three pillars close one runtime loop (Fig. 1): co-designed power electronics report real-time energy state, the synthesis layer configures the design and its algorithms, and the team executes and adapts according to its energy. Concretely:

Power system co-design. We must transition to active co-design solutions where robots proactively determine the control profiles required to automatically characterize their energy dynamics (Fig. 1a). These systems will be capable of post-deployment firmware updates to minimize energy losses in response to real-time events, akin to what biological collectives do to modulate their computation effort for survival [25, 44]. This approach will ultimately allow robots to anticipate and mitigate power failures during long-term missions.

Automated synthesis of robot designs. We need a systematic science that moves away from ad hoc design choices. Part of my current research is establishing such a framework, formulating robot designs as graphs that connect computational functions, mediating interfaces, and event-handling triggers to synthesize designs from specifications such as cost functions, reward signals or natural language (Fig. 1b). I am also developing solutions to modulate the information flow in these designs, ensuring that they can adapt in real-time.

Hardware-aware algorithms. Hardware-aware algorithms are prevalent in neuromorphic systems [14, 33], that propose computational solutions whose functionalities can be seamlessly implemented by electronic circuitry. I have started to explore neuromorphic controllers for complex decision-making that only use onboard cameras as sensors (Fig. 1c). Future work will scale these approaches to robot collectives.

REFERENCES

[1] Ahmed Abdelrahman, Matias Valdenegro-Toro, Maren Bennewitz, and Paul G Plöger. A neuromorphic approach to obstacle avoidance in robot manipulation. *The International Journal of Robotics Research*, 44(5):768–804, 2025.

[2] Jan Blumenkamp, Ajay Shankar, Matteo Bettini, Joshua Bird, and Amanda Prorok. The cambridge robomaster: An agile multi-robot research platform. In *International Symposium on Distributed Autonomous Robotic Systems*, pages 439–456. Springer, 2024.

[3] Booster Robotics. Booster T1, 2026. URL https://www.booster.tech/booster-t1/. Accessed: 2026-02-03.

[4] Charlotte Cathcart, María Santos, Shinkyu Park, and Naomi Ehrich Leonard. Proactive opinion-driven robot navigation around human movers. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4052–4058, 2023.

[5] Nick Cheney, Josh Bongard, Vytas SunSpiral, and Hod Lipson. Scalable co-optimization of morphology and control in embodied machines. *Journal of The Royal Society Interface*, 15(143):20170937, 2018.

[6] Wenceslao Shaw Cortez, Christos K Verginis, and Dimos V Dimarogonas. A distributed, event-triggered, adaptive controller for cooperative manipulation with rolling contacts. *IEEE Transactions on Robotics*, 39 (4):3120–3133, 2023.

[7] Luke de Castro, Gilhyun Ryou, Hyungseuk Ohn, and Sertac Karaman. Multi-fidelity reinforcement learning for minimum energy trajectory planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7710–7717, 2024.

[8] DJI. Camera drones selection - DJI official store. https://store.dji.com/uk/selection/camera-drones, 2026. URL https://store.dji.com/uk/selection/camera-drones. Accessed: February 3, 2026.

[9] Muhammad Umar Farooq, Amre Eizad, and Hyun-Ki Bae. Power solutions for autonomous mobile robots: A survey. *Robotics and Autonomous Systems*, 159:104285, 2023.

[10] Wojciech Giernacki, Mateusz Skwierczyński, Wojciech Witwicki, Paweł Wroński, and Piotr Kozierski. Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering. In *IEEE International Conference on Methods and Models in Automation and Robotics*, pages 37–42, 2017.

[11] Deborah M Gordon. The ecology of collective behavior in ants. *Annual Review of Entomology*, 64(1):35–50, 2019.

[12] IEEE. IEEE Xplore digital library advanced search: Multi-robot systems. https://ieeexplore.ieee.org, 2026. Search Query: ("All Metadata":"multi-robot") OR ("All Metadata":"multi robot") OR ("All Metadata":"multirobot"); Conducted on February 3, 2026.

[13] IEEE. IEEE Xplore advanced search: Multi-robot search and rescue literature. https://ieeexplore.ieee.org, 2026. Search Query: (("All Metadata":"multi-robot") OR ("All Metadata":"multi robot") OR ("All Metadata":"multirobot")) AND ("All Metadata":"search and rescue"); Conducted on February 3, 2026.

[14] Giacomo Indiveri, Bernabé Linares-Barranco, Tara Julia Hamilton, André van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih-Chii Liu, Piotr Dudek, Philipp Häfliger, Sylvie Renaud, et al. Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience*, 5:73, 2011.

[15] Claus Kadelka, Matthew Wheeler, Alan Veliz-Cuba, David Murrugarra, and Reinhard Laubenbacher. Modularity of biological systems: a link between structure and function. *Journal of the Royal Society Interface*, 20(207): 20230505, 2023.

[16] Pierre-Yves Lajoie and Giovanni Beltrame. Swarm-slam: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems. *IEEE Robotics and Automation Letters*, 9(1): 475–482, 2023.

[17] Xing Li, Oussama Zenkri, Adrian Pfisterer, and Oliver Brock. A biologically inspired design principle for building robust robotic systems. *arXiv preprint arXiv:2408.10192*, 2024.

[18] Xinfan Lin, Youngki Kim, Shankar Mohan, Jason B Siegel, and Anna G Stefanopoulou. Modeling and estimation for advanced battery management. *Annual Review of Control, Robotics, and Autonomous Systems*, 2 (1):393–426, 2019.

[19] Hou-I Liu, Marco Galindo, Hongxia Xie, Lai-Kuan Wong, Hong-Han Shuai, Yung-Hui Li, and Wen-Huang Cheng. Lightweight deep learning for resource-constrained environments: A survey. *ACM Computing Surveys*, 56(10): 1–42, 2024.

[20] Huaping Liu, Di Guo, and Angelo Cangelosi. Embodied intelligence: A synergy of morphology, action, perception and learning. *ACM Computing Surveys*, 57(7):1–36, 2025.

[21] Diego Martinez-Baselga, Eduardo Sebastián, Eduardo Montijano, Luis Riazuelo, Carlos Sagüés, and Luis Montano. Avocado: Adaptive optimal collision avoidance driven by opinion. *IEEE Transactions on Robotics*, 2025.

[22] Vito Mengers and Oliver Brock. No plan but everything under control: Robustly solving sequential tasks with dynamically composed gradient descent. In *IEEE International Conference on Robotics and Automation*, pages 90–96, 2025.

[23] David Mezey, Renaud Bastien, Yating Zheng, Neal McKee, David Stoll, Heiko Hamann, and Pawel Romanczuk. Purely vision-based collective movement of robots. *npj Robotics*, 3(1):11, 2025.

[24] Dejan Milojevic, Gioele Zardini, Miriam Elser, Andrea Censi, and Emilio Frazzoli. CODEI: Resource-efficient task-driven co-design of perception and decision making for mobile robots applied to autonomous vehicles. *IEEE Transactions on Robotics*, 41:2727–2748, 2025.

[25] Alibek Moldakozhayev and Vadim N Gladyshev. Metabolism, homeostasis, and aging. *Trends in En-*

docrinology & Metabolism, 34(3):158–169, 2023.

[26] Khai Nguyen, Sam Schoedel, Anoushka Alavilli, Brian Plancher, and Zachary Manchester. TinyMPC: Model-predictive control on resource-constrained microcontrollers. In *IEEE International Conference on Robotics and Automation*, pages 1–7, 2024.

[27] Tønnes F Nygaard, Charles P Martin, David Howard, Jim Torresen, and Kyrre Glette. Environmental adaptation of robot morphology and control through real-world evolution. *Evolutionary Computation*, 29(4):441–461, 2021.

[28] Alan Papalia, Charles Dawson, Laurentiu L Anton, Norhan Magdy Bayomi, Bianca Champenois, Jung-Hoon Cho, Levi Cai, Joseph DelPreto, Kristen Edwards, Bilha-Catherine Githinji, et al. A roadmap for climate-relevant robotics research. *arXiv preprint arXiv:2507.11623*, 2025.

[29] Nicolas Pfitzer, Eduardo Sebastián, Ajay Shankar, and Amanda Prorok. Prompting robot teams with natural language. *arXiv preprint arXiv:2509.24575*, 2025.

[30] Amanda Prorok. Extending robot minds through collective learning. *Science Robotics*, 10(106):eadv4049, 2025.

[31] José Miguel Sanz-Alcaine, Eduardo Sebastián, Iván Sanz-Gorrachategui, Carlos Bernal-Ruiz, Antonio Bono-Nuez, Milutin Pajovic, and Philip V Orlik. Online voltage prediction using gaussian process regression for fault-tolerant photovoltaic standalone applications. *Neural Computing and Applications*, 33(23):16577–16590, 2021.

[32] José Miguel Sanz-Alcaine, Eduardo Sebastián, Francisco José Perez-Cebolla, Asier Arruti, Carlos Bernal-Ruiz, and Iosu Aizpuru. Estimation of semiconductor power losses through automatic thermal modeling. *IEEE Transactions on Power Electronics*, 39(9):11086–11098, 2024.

[33] Thorben Schoepe, Ella Janotte, Moritz B Milde, Olivier JN Bertrand, Martin Egelhaaf, and Elisabetta Chicca. Finding the gap: neuromorphic motion-vision in dense environments. *Nature Communications*, 15(1):817, 2024.

[34] Eduardo Sebastián and Eduardo Montijano. Multi-robot implicit control of herds. In *IEEE International Conference on Robotics and Automation*, pages 1601–1607, 2021.

[35] Eduardo Sebastián, Eduardo Montijano, Estanis Oyarbide, Carlos Bernal, and Rubén Gálvez. Nonlinear implementable control of a dual active bridge series resonant converter. *IEEE Transactions on Industrial Electronics*, 69(5):5111–5121, 2021.

[36] Eduardo Sebastián, Eduardo Montijano, and Carlos Sagüés. Adaptive multirobot implicit control of heterogeneous herds. *IEEE Transactions on Robotics*, 38(6): 3622–3635, 2022.

[37] Eduardo Sebastián, Thai Duong, Nikolay Atanasov, Eduardo Montijano, and Carlos Sagüés. LEMURS: Learning distributed multi-robot interactions. In *IEEE International Conference on Robotics and Automation*, pages 7713–7719, 2023.

[38] Eduardo Sebastián, Eduardo Montijano, and Carlos Sagüés. ECO-DKF: Event-triggered and certifiable optimal distributed Kalman filter under unknown correlations. *IEEE Transactions on Automatic Control*, 69(4):2613–2620, 2023.

[39] Eduardo Sebastián, Eduardo Montijano, Carlos Sagüés, Mauro Franceschelli, and Andrea Gasparri. Accelerated multi-stage discrete time dynamic average consensus. *IEEE Control Systems Letters*, 7:2731–2736, 2023.

[40] Eduardo Sebastián, Mauro Franceschelli, Andrea Gasparri, Eduardo Montijano, and Carlos Sagüés. Accelerated alternating direction method of multipliers gradient tracking for distributed optimization. *IEEE Control Systems Letters*, 8:640–645, 2024.

[41] Eduardo Sebastián, Thai Duong, Nikolay Atanasov, Eduardo Montijano, and Carlos Sagüés. Physics-informed multi-agent reinforcement learning for distributed multi-robot problems. *IEEE Transactions on Robotics*, 2025.

[42] Sajad Shahsavari, Hashem Haghbayan, Antonio Miele, Eero Immonen, and Juha Plosila. A coordinated approach to control mechanical and computing resources in mobile robots. *IEEE Transactions on Robotics*, 41:347–363, 2024.

[43] Xiaochen Shang, Luo Pengwei, Xinning Wang, Jiayue Zhao, Huilin Ge, Bo Dong, and Xin Yang. Fully autonomous neuromorphic navigation and dynamic obstacle avoidance. In *Conference on Neural Information Processing Systems*, pages 1–20, 2025.

[44] Evan D Shaulson, Alan A Cohen, and Martin Picard. The brain–body energy conservation model of aging. *Nature Aging*, 4(10):1354–1371, 2024.

[45] Skydio, Inc. Skydio X10, 2026. URL https://www.skydio.com/x10. Accessed: 2026-02-03.

[46] Andrea Soltoggio, Eseoghene Ben-Iwhiwhu, Vladimir Braverman, Eric Eaton, Benjamin Epstein, Yunhao Ge, Lucy Halperin, Jonathan How, Laurent Itti, Michael A Jacobs, et al. A collective AI via lifelong learning and sharing at the edge. *Nature Machine Intelligence*, 6(3): 251–264, 2024.

[47] Francesco Stella and Josie Hughes. The science of soft robot design: A review of motivations, methods and enabling technologies. *Frontiers in Robotics and AI*, 9:1059026, 2023.

[48] Pratap Tokekar, Nikhil Karnad, and Volkan Isler. Energy-optimal trajectory planning for car-like robots. *Autonomous Robots*, 37(3):279–300, 2014.

[49] Unitree Robotics. Unitree Go2, 2026. URL https://www.unitree.com/go2. Accessed: 2026-02-03.

[50] Yuxing Wang, Zhiyu Chen, Tiantian Zhang, Qiyue Yin, Yongzhe Chang, Zhiheng Li, Liang Wang, and Xueqian Wang. Embodied co-design for rapidly evolving agents: Taxonomy, frontiers, and challenges. *arXiv preprint arXiv:2512.04770*, 2025.

[51] Stuart A West, Guy A Cooper, Melanie B Ghoul, and Ashleigh S Griffin. Ten recent insights for our understanding of cooperation. *Nature ecology & evolution*, 5

(4):419–430, 2021.

[52] Heedo Woo, Kazi RI Sanim, Keisuke Okumura, Guang Yang, Ajay Shankar, and Amanda Prorok. Sanity: An agile brushless quadrotor for multi-agent experiments. In *OpenReview: https://openreview. net/forum*, 2025.

[53] Chen Wu, Hashem Haghbayan, Abdul Malik, Antonio Miele, and Juha Plosila. Runtime energy-efficient control policy for mobile robots with computing workload and battery awareness. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 16314–16321, 2025.

[54] Liming Xin, Zhiqiu Huang, Yanying Wang, Jinlin Peng, Hao Kong, Yuehua Liu, and Yu Sun. Directed morphological evolution for fast robot structure optimization. *IEEE Transactions on Evolutionary Computation*, 2026.

[55] Yao Yao, Yuwen Chen, Liang He, and Perla Maiolino. Design optimization for bellow soft pneumatic actuators in shape-matching. In *IEEE International Conference on Soft Robotics*, pages 1–7, 2023.

[56] Sha Yi, Xueqian Bai, Adabhav Singh, Jianglong Ye, Michael T Tolley, and Xiaolong Wang. Co-design of soft gripper with neural physics. *arXiv preprint arXiv:2505.20404*, 2025.

[57] Gioele Zardini, Nicolas Lanzetti, Andrea Censi, Emilio Frazzoli, and Marco Pavone. Co-design to enable user-friendly tools to assess the impact of future mobility solutions. *IEEE Transactions on Network Science and Engineering*, 10(2):827–844, 2022.

[58] Jiliang Zhao, Wei Peng, Handing Wang, Wen Yao, and Weien Zhou. A morphological transfer-based multi-fidelity evolutionary algorithm for soft robot design. *IEEE Computational Intelligence Magazine*, 19(4):16–30, 2024.