Proceedings of the 20th Conference on Computer Science and Intelligence Systems (FedCSIS) pp. 765–770 DOI: 10.15439/2025F7678 ISSN 2300-5963 ACSIS, Vol. 43

An RL Agent to Find Minimum Energy in a Tensegrity Representing a Cell

Mustafa Shah
Department of Electronics,
Quaid-e-Azam University,
Islamabad, Pakistan
mustafamohmand59@gmail.com

Arsenio Cutolo
Dept. of Structures for Engineering and Architecture
University of Napoli Federico II
Naples, Italy
arsenio.cutolo@unina.it

Muddasar Naeem
Università Giustino Fortunato
Benevento Italy
m.naeem@unifortunato.eu

Muhammad Waris
Department of Electronics,
Quaid-e-Azam University,
Islamabad, Pakistan
mwaris.22411012@ele.qau.edu.pk

Musarat Abbas
Department of Electronics,
Quaid-e-Azam University,
Islamabad, Pakistan
mabbas@qau.edu.pk

***Abstract*—Understanding the mechanical behavior of cells is a complex challenge at the crossroads of physics, biology, and engineering. The cytoskeleton which is a dynamic network of filaments which helps cells maintain shape, move, and respond to their environment. Tensegrity structures, made of interconnected tensile and compressive elements, offer a compelling way to model these internal forces. In this work, we use Reinforcement Learning (RL) to simulate and optimize cellular mechanics. We propose an RL framework where an agent learns to minimize the total mechanical energy of tensegrity-based cell models by adjusting node positions. We consider diverse shapes from simple shapes like lines and triangles, to more complex shapes like cell-like geometries. Our approach shows that RL can effectively model mechanical adaptations in cells and opens the door to intelligent, bio-inspired simulations. This work bridges biophysics, AI, and structural mechanics, offering new ways to predict and understand how cells respond to mechanical stress.**

***Index Terms*—Reinforcement Learning, Tensegrity, Cell Mechanics, Structural Optimization, Deep Learning, PPO, Cytoskeleton**

I. Introduction

CELLS are not static entities; they are dynamic, mechanically active systems that continuously adapt their internal architecture in response to various environmental and biological cues [14]. One of the most compelling frameworks for modeling the mechanical integrity of cells is the tensegrity model, which represents the cytoskeletal network as a balance of tensile and compressive forces [1], [4]. Tensegrity—short for "tensional integrity"—describes a structural system composed of isolated struts (under compression) suspended within a web of tensile elements (like cables or filaments). This principle reflects the way real cells maintain their shape, transmit forces, and react to mechanical stress.

The tensegrity-based approach allows one to capture the interplay between cytoskeletal elements such as actin filaments, microtubules, and intermediate filaments, providing a robust platform to simulate how cells respond to external mechanical stimuli and internal biochemical signals [3]. Beyond biology, tensegrity models have also found significant applications in the design and development of soft robotic systems and bioinspired structural materials that mimic cellular mechanics and adaptability [15].

With the rapid advancement of artificial intelligence and machine learning, particularly in the realm of reinforcement learning (RL), new opportunities have emerged to understand and control dynamic systems in ways that were previously unfeasible [8]. RL algorithms allow agents to interact with complex environments and learn control strategies through trial and error [13].

In this work, we propose a novel integration of RL with tensegrity physics to simulate a biologically inspired cell model. An RL agent is tasked with iteratively adjusting the positions of the nodes within a tensegrity structure to achieve a configuration with minimal mechanical energy, a proxy of structural equilibrium. Previous work in robotics has demonstrated the efficacy of such approaches in enabling tensegrity-based systems to learn locomotion and adaptive behaviors [6], [17]. By extending these concepts to cellular modeling, we aim to bridge the gap between mechanobiology and intelligent control systems.

The rest of the paper is organised as follow: Section II presents related work where we review the exisitng work in cell mechanics as well as use of RL to solve various problems. The section III is about our system model and methodology were we present in details our approach and tensegrity structure visualizations using RL. We outlined the results of experiments in section IV, followed by a discussion in section V and conclusion in section VI.

II. Related Work

In this section, we report the relevant literature to highlight the limitation of existing work and contribution of our work.

The work in [18] proposes an interpretable machine learning framework of cell mechanics from protein images and neural networks (NN) for prediction of traction forces from a single focal adhesion protein field. Agnostic and physics-constrained approaches learn interpretable rules for prediction. A DCell is proposed in [20] that is a visible NN embedded in the hierarchical structure of 2,526 subsystems and encompasses an eukaryotic cell. DCell is able to simulates cellular growth with reasonable accuracy after training on millions of genotypes. The proposed setup provides a foundation to decode the drug resistance, genetics of disease, and synthetic life.

An interesting work is done in [22] to develop the safe RL approaches specially for environments where non rewarded states are also important. Authors considered different scenario to validate their approaches and use of safe RL methods could have useful contribution in cell mechanics. RL methods are also used in important medical issues to provide assistance to patients with disabilities [19], [21]. These works demonstrate a useful application of RL algorithms but the such formulation of tensegrity representing a cell is work to be done.

Moreover, Deep RL (DRL) is utilized in [9] to infer collective cell behaviours and cell–cell interactions in tissue morphogenesis from 3D time-lapse images. Hierarchical DRL is applied to investiage cell migrations from the images with an ubiquitous nuclear label. The hierarchical DRL method HDRL reveals a modular, multiphase organization of cell movement to Caenorhabditis elegans embryogenesis. A hybrid RL model is proposed in [12] to manage process control efficiently. A probabilistic knowledge graph mdoel is developed characterizing the science and risk-based understanding of quantifying inherent stochasticity and biomanufacturing process mechanisms.

In our work, we investigate how reinforcement learning can be applied to better understand the ways in which biological cells adapt their shapes and internal structures in response to mechanical forces. The proposed work formulate a tensegrity model representing a cell into an RL setup and then design a RL agent using proximal policy optimization to find minimum energy in a tensegrity representing a cell.

III. Methodology

This section presents the methodology used in our work. The first step is to model the cell as a tensegrity structure—a network of compressive elements (rods) and tensile elements (cables)—to capture how real cells maintain mechanical balance through tension and compression [1].

To explore this model computationally, we employ a Reinforcement Learning (RL) agent that learns to adjust node positions in the tensegrity structure to reduce the total mechanical energy, guiding it toward a stable, low-energy configuration. We use the Proximal Policy Optimization (PPO) algorithm due to its robustness in continuous control tasks [2].

By integrating biophysics, artificial intelligence, and structural mechanics, this approach contributes to the understanding of how cellular behavior emerges from physical constraints. Such methods also support the design of soft robotic systems that mimic cellular mechanics [5].

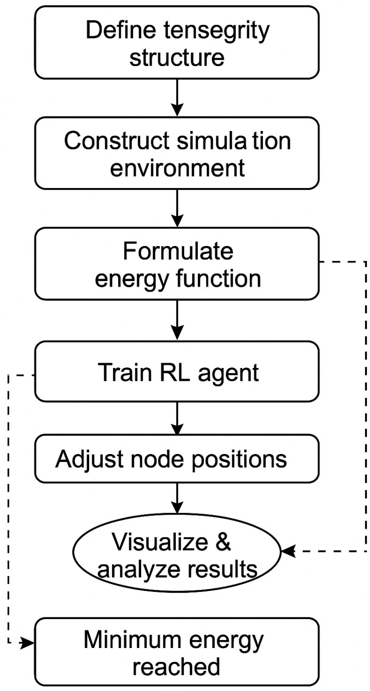

Fig. 1: Flowchart outlining the proposed approach.

The mechanical energy of the tensegrity structure is computed based on edge deformations relative to rest length:

$$E = \sum_{(i,j)\in \text{edges}} k_{ij} \cdot (||x_i - x_j|| - l_{ij}^0)^2$$

where x_i, x_j are 2D positions of nodes, l_{ij}^0 is the rest length, and k_{ij} is the stiffness constant (set to 1). This energy reflects internal stress.

1) Mechanical Energy Decomposition: To distinguish tension and compression, we decompose the total energy:

$$E_{\text{total}} = \sum_{(i,j)\in\mathcal{C}} k_{ij}^{(c)} \left(\|x_i - x_j\| - l_{ij}^0\right)^2 + \sum_{(i,j)\in\mathcal{T}} k_{ij}^{(t)} \left(\|x_i - x_j\| - l_{ij}^0\right)^2 \quad (1)$$

Here, $\mathcal{C}$ and $\mathcal{T}$ are sets of compressive and tensile elements, with different stiffness coefficients $k_{ij}^{(c)}$, $k_{ij}^{(t)}$, respectively.

2) Gradient of Energy with Respect to Node Position: Although PPO is model-free, internal forces can be approximated by the energy gradient:

$$\frac{\partial E}{\partial x_i} = \sum_{j\in\mathcal{N}(i)} 2k_{ij}(\|x_i - x_j\| - l_{ij}^0) \cdot \frac{(x_i - x_j)}{\|x_i - x_j\|} \quad (2)$$

This gradient represents the net force acting on node i, offering a physical interpretation of movement.

3) Proximal Policy Optimization (PPO) Loss: We use PPO's clipped surrogate loss [2]:

$$\mathcal{L}^{\text{PPO}}(\theta) = \mathbb{E}_t \left[\min \left(r_t(\theta)\hat{A}_t, \text{clip}\left(r_t(\theta), 1-\epsilon, 1+\epsilon\right)\hat{A}_t \right) \right] \tag{3}$$

with

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$$

where $\hat{A}_t$ is the advantage estimate. This formulation ensures stable policy updates.

4) Energy Convergence Metric: To evaluate learning stability, we compute average energy change over the last N steps:

$$\Delta E = \frac{1}{N} \sum_{t=T-N+1}^{T} |E_t - E_{t-1}| \tag{4}$$

A small ΔE signals convergence to a stable, low-energy configuration.

5) RL Formulation:

- **State Space:** A flattened vector of 2D node coordinates representing the structure.
- **Action Space:** Small shifts $(\Delta x, \Delta y)$ applied to movable nodes.
- **Reward Function:** Negative of total mechanical energy.
- **Termination:** Fixed-length episodes without early stopping.

6) Implementation Steps:

- Model the cell as a tensegrity structure of nodes and rods/cables [3].
- Develop a Python simulation environment to update node positions.
- Define an energy function to evaluate mechanical stress.
- Use PPO (via Stable-Baselines3) to train the agent [2].
- Train the agent to iteratively reduce system energy.
- Visualize learning and adaptations of the structure.
- Analyze biological relevance of learned behaviors [5].
- Explore applications in mechanobiology and soft robotics [15].

A. Reproducibility and Future Extensions

Each structure is defined by nodes and edges (cables/rods). A custom Python environment updates positions using actions from the agent. The state is a flattened list of node coordinates, while actions apply perturbations. The reward is based on internal potential energy. All components are modular and can be adapted for future extensions, including 3D structures and additional biophysical constraints.

B. Tensegrity Structures Visualization

We consider diverse tensegrity structures, from simple to complex, in our work as shown in Figure 2. Each configuration is simulated and visualized using Python, enabling us to observe how reinforcement learning (RL) guides the system toward more stable, low-energy states. These structures ranging from basic geometries to biologically inspired shapes demonstrate the adaptability of RL in modeling cellular mechanics [6], [7].

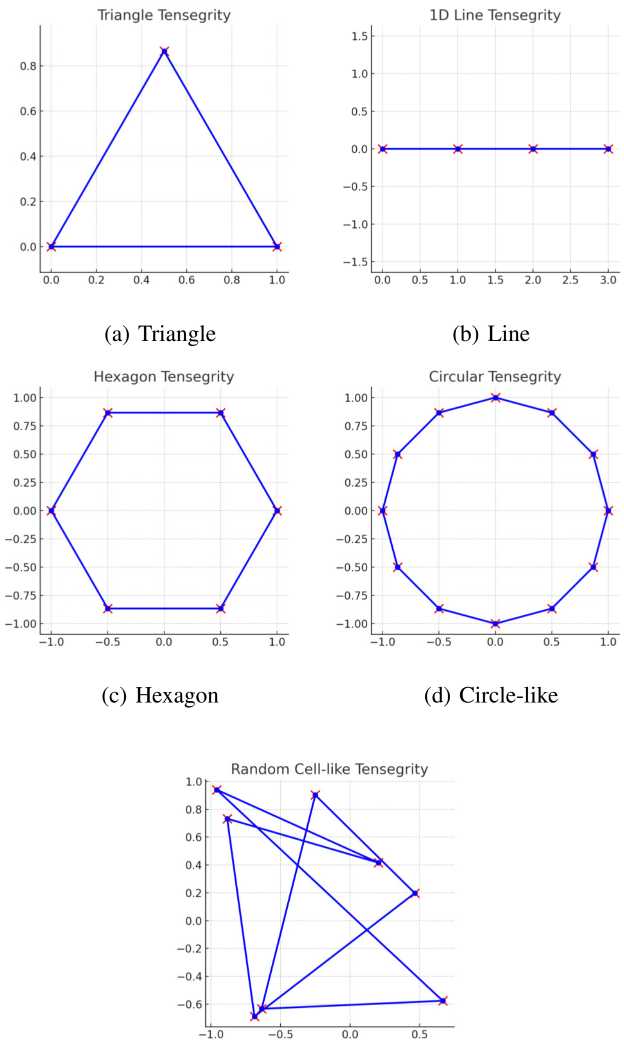

(a) Triangle (b) Line (c) Hexagon (d) Circle-like (e) Irregular Cell-like

Fig. 2: Various tensegrity structures used in the study.

Each of the structures illustrated in Figure 2 represents a unique scenario in our simulations. Figure 2a shows a simple triangular tensegrity, serving as a basic case to validate the RL agent's energy minimization capability in a constrained geometry. Figure 2b presents a linear configuration, ideal for examining the agent's behavior in elongated systems. The symmetric hexagonal form in Figure 2c introduces more complexity, resembling biological structures. Figure 2d illustrates a radially balanced, circular tensegrity, while Figure 2e presents an irregular, cell-like morphology—a step toward realistic cellular modeling.

Across all cases, the RL agent effectively minimized system energy, reinforcing the concept that even simple tensegrity models can provide meaningful insights into cellular biomechanics.

C. Towards Realistic Cell-like Tensegrity Structures

We demonstrate a biologically inspired tensegrity system that reflects the intricate architecture of real cells, in Figure 3. This advanced model incorporates elements resembling cytoskeletal components such as actin filaments and microtubules, which are critical for maintaining cell shape and enabling mechanical responsiveness. Such structures have been central to tensegrity-based modeling frameworks in cellular biomechanics [4], [5]. By simulating these complex geometries, we move closer to capturing the adaptive, dynamic nature of living cells.

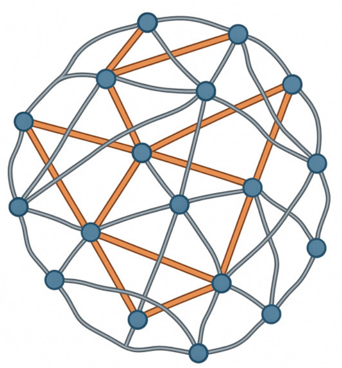

Fig. 3: A biologically inspired tensegrity model representing a complex cell.

IV. Results

This section presents the experimental results that we performed while considering different structures.

Our experiments began with basic triangular tensegrity structures as shown in Figure 2a, where the RL agent consistently achieved noticeable reductions in total mechanical energy over training episodes. These early successes demonstrated the agent's capacity to learn stable configurations, even in minimal geometric setups.

As we progressed to more complex shapes such as hexagonal (Figure 2c) and circular configurations (Figure 2d), the RL agent adapted well, maintaining structural coherence while reducing energy levels. Notably, the hexagonal structures, which resemble symmetrical cell geometries, converged faster toward low-energy states. The circular models, simulating soft boundary conditions like those found in real cell membranes, further confirmed the agent's ability to generalize across geometries.

Visual analysis of the final tensegrity forms and energy plots validated the model's effectiveness. These findings are consistent with earlier research where RL-based policies achieved adaptive control and energy-efficient locomotion in soft tensegrity robots [7], [11].

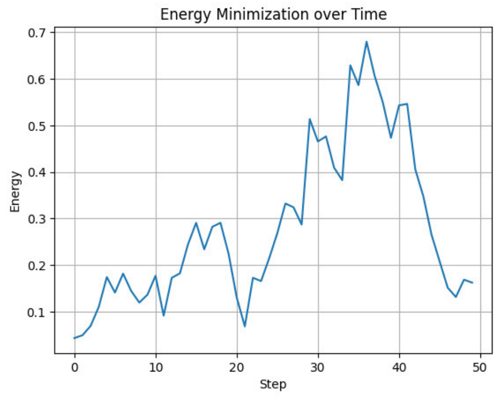

(a) Triangle tensegrity

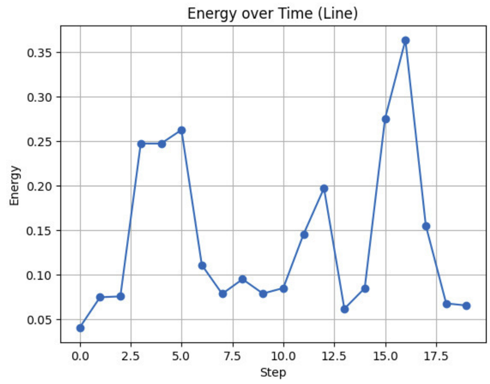

(b) Line tensegrity

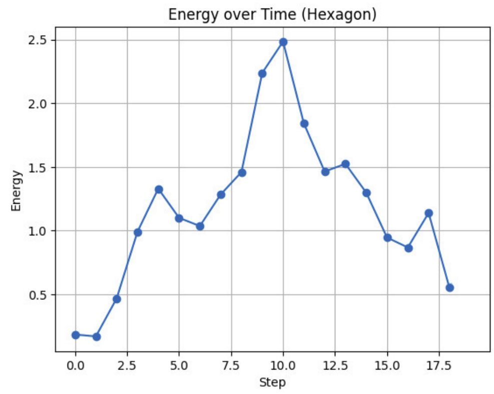

(c) Hexagon tensegrity

Fig. 4: Energy of the tensegrity structures over training steps. (a) Triangle tensegrity (b) Line tensegrity (c) Hexagon tensegrity.

For tensegrity structures to train the RL agent, we employed the Proximal Policy Optimization (PPO) algorithm and used a reward function based on the system's mechanical energy. Specifically, the agent receives a negative reward equal to the sum of squared differences between the current and rest lengths of all links, effectively encouraging configurations that minimize internal tension and compression. The environment state is defined as a flat vector of 2D coordinates representing the positions of all nodes. This simple representation provides the agent with sufficient spatial information to act. During early training, exploration is naturally handled through PPO's stochastic policy updates, allowing the agent to try unstable configurations and gradually converge toward energy-efficient ones.

Each experiment was trained over 100,000 timesteps, typically spanning 500 to 1,000 episodes depending on episode length. This duration provided the agent with enough interaction to consistently learn optimal or near-optimal configurations. For example, in the triangular and hexagonal tensegrities, the policy reliably minimized energy within the first 60,000 timesteps, while irregular cell-like shapes required more training to stabilize. We observed that once the energy starts decreasing steadily, subsequent steps fine-tune the positions rather than cause drastic shifts. Preliminary robustness checks indicate the learned policies remain effective when tested on slightly perturbed initial states.

While our current results are mostly qualitative, they offer clear and reproducible patterns of energy reduction across varied structures. In future work, we plan to include statistical metrics such as average final energy, convergence rate, and baseline comparisons with classical physics-based solvers. We are also extending our approach to 3D tensegrity models and exploring more expressive RL methods like Soft Actor-Critic (SAC) to improve learning in high-dimensional, nonlinear cellular geometries.

V. Discussion and Future Work

This study has demonstrated the feasibility of using reinforcement learning to manipulate tensegrity-based cell models toward stable and energy-efficient states. It provides a promising step in merging biophysical modeling with artificial intelligence, especially for systems that rely on structural tension and compression, like real cells. Despite its success in 2D structures, our framework opens the door to deeper biological relevance through future improvements as explained below:

- **Adding biological constraints**, such as anchor points representing extracellular matrix contacts or internal pressure mimicking cytoplasmic resistance [4].
- **Scaling to 3D tensegrity systems**, which can better replicate the spatial complexity of cytoskeletal networks found in living cells [10].
- **Incorporating viscoelasticity**, allowing simulation of how cells respond over time to mechanical stimuli and not just at equilibrium.
- **Integrating with real data**, such as microscopy-derived cell geometries, to improve biological validity and potentially guide experimental modeling [5].

Finally, we can summarize that as artificial intelligence continues to make significant contributions in the medical and healthcare domains, its integration into biological modeling is becoming increasingly relevant. Studies have shown how AI-driven approaches enhance diagnostic precision in areas such as skin cancer and COVID-19 detection [23], as well as in risk management [22]. In a similar spirit, our work demonstrates the potential of reinforcement learning for understanding complex biomechanical behavior at the cellular level. This further supports the growing trend of leveraging AI not only for diagnosis but also for predictive and mechanistic modeling of biological systems.

These enhancements would bring the relevant research community closer to a useful tool that can not only simulate the structure but also predict and interpret cellular mechanical behavior in health and disease.

In the future work, we will extend our work to more complex topologies and we will also consider including external constraints like cell-substrate interaction or intracellular pressure.

VI. Conclusion

In this work, we have presented a novel integration of tensegrity mechanics with reinforcement learning to explore how AI methods, specifically RL agent can autonomously discover stable, low-energy configurations of cell-like structures. By learning control strategies that minimize mechanical energy, the agent mimics cellular adaptation behaviors in silico. This framework lays the groundwork for more intelligent models of mechanobiology—capable of interpreting, simulating, and even designing cellular responses. With further development, it may contribute to advances in synthetic biology, bio-inspired robotics, and our theoretical understanding of cellular mechanics.

Acknowledgment

This research has been partially supported by the European Union - Next Generation EU through the Project of National Relevance ”Innovative mathematical modeling for cell mechanics: global approach from micro-scale models to experimental validation integrated by reinforcement learning”, financed by European Union-Next-GenerationEU-National Recovery and Resilience Plan-NRRP-M4C1-I 1.1, CALL PRIN 2022 PNRR D.D. 1409 14-09-2022—(Project code P2022MXCJ2, CUP F53D23010080001) granted by the Italian MUR.

References

[1] D. E. Ingber, “Cellular tensegrity: defining new rules of biological design that govern the cytoskeleton,” *Journal of Cell Science*, vol. 104, pp. 613–627, 1993.

[2] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” arXiv preprint arXiv:1707.06347, 2017.

[3] R. E. Skelton and M. C. de Oliveira, *Tensegrity Systems*, Springer, 2009.

[4] D. E. Ingber, “Tensegrity I. Cell structure and hierarchical systems biology,” *Journal of Cell Science*, vol. 116, no. 7, pp. 1157–1173, 2003.

[5] Y. Chen, X. Wang, L. Yang, et al., “Machine Learning for Mechanobiology: From Cells to Tissues,” *Annual Review of Biomedical Engineering*, vol. 25, pp. 377–406, 2023.

[6] C. Paul and M. Spröwitz, “Reinforcement learning for tensegrity robot locomotion,” *Frontiers in Robotics and AI*, vol. 5, p. 77, 2018.

[7] K. Caluwaerts, J. Despraz, A. Işçen, A. Sabelhaus, V. SunSpiral, and R. D. Beer, “Design and control of compliant tensegrity robots through simulation and hardware validation,” *Journal of The Royal Society Interface*, vol. 11, no. 98, 2014.

[8] Naeem, M., Coronato, A., Ullah, Z., Bashir, S. and Paragliola, G., 2022. Optimal user scheduling in multi antenna system using multi agent reinforcement learning. Sensors, 22(21), p.8278.

[9] Wang, Z., Xu, Y., Wang, D., Yang, J. and Bao, Z., 2022. Hierarchical deep reinforcement learning reveals a modular mechanism of cell movement. Nature machine intelligence, 4(1), pp.73-83.

[10] V. SunSpiral, A. Agogino, D. Atkinson, and A. Moored, "Tensegrity based probes for planetary exploration: Entry, descent and landing (EDL) and surface mobility analysis," *NASA Ames Research Center*, 2013.

[11] M. Ahn, J. Lee, S. Lee, and Y. Lee, "Policy gradient reinforcement learning for locomotion control of tensegrity robots," *International Conference on Ubiquitous Robots*, 2019.

[12] Zheng, H., Xie, W., Wang, K. and Li, Z., 2022. Opportunities of Hybrid Model-based Reinforcement Learning for Cell Therapy Manufacturing Process Control. arXiv preprint arXiv:2201.03116.

[13] D. Pathak, P. Agrawal, A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," *Proceedings of the 34th International Conference on Machine Learning*, 2017.

[14] Naeem, M., Fiorino, M., Addabbo, P. and Coronato, A., 2024. Integrating Artificial Intelligence Techniques in Cell Mechanics. Annals of Computer Science and Information Systems, 41, pp.111-116.

[15] M. Pagitz, M. Pellegrino, and A. Daynes, "Bio-inspired tensegrity soft robots," *Bioinspiration & Biomimetics*, vol. 7, no. 4, 2012.

[16] Fiorino, M., Naeem, M., Ciampi, M. and Coronato, A., 2024. Defining a metric-driven approach for learning hazardous situations. Technologies, 12(7), p.103.

[17] J. Kim, Y. Kim, and J. Ryu, "Bio-inspired reinforcement learning control for soft tensegrity robot locomotion," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3937–3944, 2022.

[18] Schmitt, M.S., Colen, J., Sala, S., Devany, J., Seetharaman, S., Caillier, A., Gardel, M.L., Oakes, P.W. and Vitelli, V., 2024. Machine learning interpretable models of cell mechanics from protein images. Cell, 187(2), pp.481-494.

[19] Ismail, A., Naeem, M., Khalid, U.B. and Abbas, M., 2025. Improving adherence to medication in an intelligent environment using reinforcement learning. Journal of Reliable Intelligent Environments, 11(1), pp.1-10.

[20] Ma, J., Yu, M.K., Fong, S., Ono, K., Sage, E., Demchak, B., Sharan, R. and Ideker, T., 2018. Using deep learning to model the hierarchical structure and function of a cell. Nature methods, 15(4), pp.290-298.

[21] Ismail, A., Naeem, M., Syed, M.H., Abbas, M. and Coronato, A., 2024. Advancing Patient Care with an Intelligent and Personalized Medication Engagement System. Information, 15(10), p.609.

[22] Fiorino, M., Naeem, M., Ciampi, M. and Coronato, A., 2024. Defining a metric-driven approach for learning hazardous situations. Technologies, 12(7), p.103.

[23] Qayyum H., Rizvi S.T.H., Naeem M., Khalid U.b., Abbas M., and Coronato A., "Enhancing Diagnostic Accuracy for Skin Cancer and COVID-19 Detection: A Comparative Study Using a Stacked Ensemble Method," *Technologies*, vol. 12, no. 9, p. 142, 2024. doi: https://doi.org/10.3390/technologies1209014210.3390/technologies12090142