nature portfolio

Peer Review File

Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery

Corresponding Author: Professor Sebastian Risi

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The manuscript titled ' Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery ' presents an innovative study: a biologically inspired, fully decentralized physical cellular system capable of distributed morphological recognition, fault tolerance, and even simulated self-repair. Before publication, the manuscript should be revised as follows
1) While the tested shapes (airplane, chair, etc.) are diverse but relatively simple. The system's recognition and repair performance should be assessed on more complex, intricate structures or those with internal cavities (e.g., human models, machinery). Would the current network scale and training strategy suffice for such complexity, or would adjustments be necessary?
2) The manuscript introduces a novel decentralized method (Neural Cellular Automaton) but lacks a comparative analysis. To properly demonstrate its advantages, it is essential to include benchmarks against traditional or simpler decentralized algorithms. Quantitative comparisons in terms of accuracy, convergence speed, and robustness are needed to validate the method's contributions.
3) The Conclusion could be strengthened in two ways. First, the authors should specify potential application scenarios for this technology. Second, a discussion of the gaps between the system and real biology (e.g., in energy acquisition, cellular differentiation, or scale invariance) would better outline future research challenges and directions.
4) The results section has an inconsistent narrative flow. The authors describe Figures 2A-D, jump to Figure 3, and then return to Figures 2E-H. We recommend presenting all of Figure 2 (A-H) together to first fully describe the classification capabilities, followed by the robustness analysis in Figure 3. This will significantly improve the logical coherence.

Reviewer #2

(Remarks to the Author)
This work presents the first 3D hardware realization of decentralized morphogenesis and self-recognition. Its most significant contribution is the physical demonstration of a collective of nearly 200 independent modules achieving consensus on their 3D global shape (e.g., “plane,” “table,” “boat”) without centralized control or global positioning, relying solely on local communication. The 3D convolutional design, featuring a cross-shaped kernel, is well-motivated and draws inspiration from biological systems. The use of PCBs as structural components is both elegant and cost-effective.

However, the experimental setup and evaluation of results suffer from notable limitations. While the physical hardware implementation is constrained to a lower resolution (15×15×15) due to practical considerations, all experiments on robustness and damage recovery were conducted entirely in simulation. From a computer graphics research perspective, such experiments could—and arguably should—be performed at higher resolutions (e.g., 64×64×64 voxels) to yield more meaningful and generalizable insights. The current results do not provide sufficient evidence that the NCA-based approach will perform similarly at higher resolutions, limiting the scalability claims.

Furthermore, the self-recognition and damage detection algorithms are implemented as separate processes, despite both being based on the same NCA framework. Notably, damage detection requires conditioning the NCA on a given shape class. If biological systems can simultaneously exhibit both capabilities, it is reasonable to expect that the neural cells in this system should also be capable of concurrent self-recognition and damage detection. Since the experiments already demonstrate robust classification under 5–15% fault conditions, extending the framework to detect and localize these faults within the same process would be a more significant and coherent step toward true decentralized self-organization.

Reproducibility is also hindered by insufficient detail in several key areas:

1. Cell State Initialization at t=0: The manuscript does not specify the initial values for the 28 state channels (hidden, logits, alpha), which are critical for reproducibility.
2. Inter-cell Communication Content: The manuscript lacks a clear description of what “information” is aggregated from neighboring cells. Do cells exchange the full 28D state vector, a subset of it, or some processed perceptual features?
3. Consistency in Damage Identification: The algorithm does not explain how consistency is enforced among different cells during damage identification. For example, if two normal cells, A and B, are both adjacent to a region C, and cell A identifies C as damaged while cell B does not, what mechanism resolves this conflict?

These omissions significantly impede faithful re-implementation and validation of the proposed system.

In summary, while the physical demonstration of decentralized shape classification is impressive, the work remains preliminary. The reliance on low-resolution simulations, the separation of recognition and fault detection, and the lack of implementation details prevent a full assessment of scalability and robustness. Given these unresolved core issues, I recommend a major revision. Essential revisions should include:
1. Experiments conducted at higher resolutions;
2. A unified algorithm enabling simultaneous shape recognition and fault detection;
3. Clear and complete specifications for initialization, inter-cell communication, and conflict resolution mechanisms.

Reviewer #3

(Remarks to the Author)
Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery

This paper introduces the concept of a ‘smart brick’, which contains its own local sensing, compute – predominantly used to run a local neural network - and communication, making each brick an autonomous agent in its own right. The bricks can be combined into a wide variety of morphologies.

The authors run a good variety and depth of experimentation, characterising emergent system properties including shape classification, fault tolerance, and shape repair. Compared to the state of the art, these bricks implement full decentralisation, with experimentation demonstrating capabilities that exceed state of the art.

Presentation is pretty good, it feels like a Transactions on Graphics paper. Figures are useful, well-captioned, and clearly communicate key features of the proposed approach and related experimentation.

The authors should comment on physical scalability – they simulate hundreds of blocks, and systems like this make most sense with a large number of blocks, e.g. into the hundreds. The blocks themselves are made out of PCBs and seem to require manual assembly e.g., to cable tie the ‘lid’ on. How long does it take to create one block, and how much manual effort is required? How expensive is a block? It seems like fabrication may be a bottleneck to any practical realisation.

A big issue here is the practical application. The motivation given seems to be predominantly around ‘modular robotics’ and ‘cellular repair’, but there is no actual use case provided. Are the blocks to make up a robot? Are they to repair a larger structure? What structure could benefit from decentralised self-repair and self-classification, and be essentially made of PCBs at the given scale? Practical applicaiton seems to be well out of reach of the current technology, which should be highlighted more forthrightly in the conclusions if I am correct. I am also interested in how small such bricks could be made – do the authors have any comments on this?

Actuation is a major issue. I recognise that the authors discuss this in the ‘limitations’ section, which is good. However, most of the capabilities demonstrated in the simulation cannot currently be transferred to the real agents. This compounds my above issue with practical application and fabrication scalability and suggests that the physical block systems of reasonable size and complexity will be prohibitively difficult to create. It would also be helpful to know how sturdy and robust these bricks are – can the authors perform some testing with significant numbers of cycles of assembly/disassembly, even for two bricks being repeatedly connected and disconnected, to get some idea of robustness and implications for reusability?

The core limitation with the damage detection is that it only works on a specific instance of a given class, rather than across a class. It would be good to see what happens when this is done per class, and what the failure modes are – e.g., what makes a table repair itself into a chair, and can chains of local activations be somehow visualised? This would be a very interesting study to see, as you could track the emergent behaviour over time; e.g., does a broken table start to repair itself as a table, then switch to a chair? Under what conditions does it switch between classes? It must also be said that this combines with the rest of the limitations in this study to somewhat reduce the impact of the paper overall.

Despite this, there are some solid contributions. Self-repair and fully distributed classification (which even somewhat holds across different scales and is therefore scalable!) are fundamental features of any successful future modular robotic/structural system. The authors draw on some biological inspiration, but it is unclear to me if biological self-repair works anything like the mechanisms proposed in the paper. If this biological inspiration has merit, it should be expanded on and concretely demonstrated. If not, it should be rewritten to appropriately represent the extent to which the proposed approach is bio-inspired.

Minor typo “no more damaged was detected.”

Overall I’m torn on this one. It’s certainly interesting, well written, and well presented. It has significant limitations that lessen impact and, in particular, stop the core potential contributions being practically realised. Solving these limitations would involve significant effort in hardware development but would improve the article and make it a very high quality Nature Comms article. I'm hesitant to advocate for acceptance without the above-mentioned additions being strongly considered by the authors.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors have addressed my comments

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)
The authors have adequately addressed the concerns raised in the previous round. The revised manuscript now includes higher-resolution experiments (32^3 and 64^3), a joint self-recognition and damage detection model, a WL baseline comparison, and implementation details. These additions satisfactorily resolve my earlier concerns about scalability, concurrency, and reproducibility.

The manuscript in its current form presents a clearer and more complete case for its contributions. I have only a few minor suggestions regarding presentation:

1. Narrative flow: The WL baseline comparison and high-resolution results currently appear in the Discussion. I recommend moving them into the Results section, so the narrative flows more smoothly.

2. Figure 2H: On page 17, the text refers to “Figure 2H,” but subpanel H within Figure 2 appears to be missing or has not been updated. Please verify.

3. Consistency in supplementary figures: In Figures S5–S7, the color for the same algorithm varies across panels. A consistent color assignment for each method would make comparisons easier.

Overall, the revised work is substantially stronger. The remaining issues are minor and pertain only to presentation. Once addressed, I believe the manuscript is suitable for publication.

(Remarks on code availability)
The .npy and .zip data files at the root level would be better placed in a dedicated subfolder, and adding a .gitignore to exclude intermediate files such as .vscode/ and __pycache__/ would keep the repository cleaner.

Reviewer #3

(Remarks to the Author)
The manuscript has been improved and many of the limitations have been either addressed or more specifically called out. The justification vs. biology is now fairer.

I appreciate the scaling baseline, however the computational cost of scaling in simulation must be explicitly called out. This needs adding to the manuscript.

Otherwise I am happy that the paper is now improved. Given the editor is happy with the fundamental limitations of the approach and questionable real-world applicability that I previously mentioned, this paper - once computational cost is added - is good enough to be accepted.

(Remarks on code availability)

We thank the editor and reviewers for their careful and constructive evaluation of our manuscript, *“Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery.”* The feedback was highly valuable and led to substantial revisions throughout the paper.

In response to these concerns, we significantly strengthened the manuscript along the key axes raised during review: scalability, comparative evaluation, methodological clarity, practical scope, and biological framing. In particular:

- **We added new higher-resolution experiments at 32×32×32 and 64×64×64, showing that the decentralized NCA framework extends beyond the original hardware-constrained regime to larger, denser, and more intricate voxelized structures when sufficient model capacity is provided.**
- **We introduced a joint model for simultaneous self-recognition and damage detection, addressing some reviewers’ concern that these capabilities should**

emerge together within a single decentralized process. We also clarified the complementary role of the shape-conditioned recovery model in resolving ambiguity under severe damage.

- **We added a decentralized baseline comparison based on a distributed Weisfeiler–Lehman approach, including quantitative comparisons in classification accuracy, robustness, and convergence properties, to better position the contribution of the NCA framework.**
- **We clarified several methodological details important for reproducibility, including cell-state initialization, the full 28-dimensional inter-cell communication state, and the local voting mechanism used to resolve damage-direction predictions.**
- **We expanded the Discussion of physical scalability and practical scope, including fabrication time, per-module cost, miniaturization prospects, and a clearer distinction between near-term applications and longer-term visions requiring substantial hardware advances in actuation and self-reconfiguration.**
- **We revised the biological framing throughout the manuscript to make clear that the work is inspired by high-level principles of decentralized local interaction and emergent organization, rather than intended as a mechanistic model of biological self-repair.**

Overall, we believe the revised manuscript is substantially stronger and now presents a clearer and more complete case for its central contribution: a physical 3D demonstration that decentralized shape classification and damage-aware behavior can emerge from identical modules using only local communication and distributed control. In this sense, we view the work as an important stepping stone toward future modular robotic and adaptive material systems.

REVIEWER COMMENTS

Reviewer #1 (Remarks to the Author):

The manuscript titled ' Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery ' presents an innovative study: a biologically inspired, fully decentralized physical cellular system capable of distributed morphological recognition, fault tolerance, and even simulated self-repair. Before publication, the manuscript should be revised as follows
1) While the tested shapes (airplane, chair, etc.) are diverse but relatively simple. The system's recognition and repair performance should be assessed on more complex, intricate structures or those with internal cavities (e.g., human models, machinery). Would the current network scale and training strategy suffice for such complexity, or would adjustments be necessary?

We thank the reviewer for this important comment. We agree that evaluating the method only in the low-resolution regime would provide limited evidence for scalability. To address this, we expanded the manuscript with additional experiments at higher resolutions, evaluating the same approach at 32 × 32 × 32 and 64 × 64 × 64. These tests

go substantially beyond the hardware-constrained regime and include much larger and more intricate voxelized structures with over 18,000 cubes (Figure 6).

The new results show that the method remains effective at higher resolutions, especially when model capacity is increased. At 15 × 15 × 15, accuracies for hidden channel sizes of 20, 48, 96, and 128 are 0.71, 0.83, 0.98, and 0.98, respectively. At 32 × 32 × 32, the corresponding accuracies are 0.40, 0.54, 0.87, and 0.96, and at 64 × 64 × 64 they are 0.38, 0.53, 0.87, and 0.96. This shows that higher-capacity models retain strong performance even as resolution increases, while smaller hidden-state sizes degrade more substantially.

We also now report training efficiency across resolutions. The model exceeds 90% accuracy after 108, 517, and 611 training steps for the 15^3, 32^3, and 64^3 settings, respectively. Thus, while larger and more complex grids require more training, the approach continues to learn efficiently at higher resolutions.

These experiments strengthen the scalability evidence in two ways. First, they show that our NCA framework is not restricted to the original low-resolution setting. Second, they indicate that scaling to larger and more complex structures is feasible, but benefits from increased hidden-state size and thus increased model capacity. We have revised the manuscript accordingly.

2) The manuscript introduces a novel decentralized method (Neural Cellular Automaton) but lacks a comparative analysis. To properly demonstrate its advantages, it is essential to include benchmarks against traditional or simpler decentralized algorithms. Quantitative comparisons in terms of accuracy, convergence speed, and robustness are needed to validate the method's contributions.

We agree that the contribution of the proposed NCA framework should be evaluated against a simpler decentralized alternative. In the revised manuscript, we have therefore added a comparative benchmark based on a distributed Voxel Graph Weisfeiler–Lehman (WL) baseline. This comparison now includes quantitative results for classification accuracy, distributed convergence speed, and robustness under perturbation.

The added results show that the two approaches achieve comparable classification performance: the WL baseline reaches 100% classification accuracy, compared with 98.97% for the NCA. We also now report robustness under damaged-cell experiments aggregated across all damaged shapes and across error rates of 5%, 10%, and 15%. In this setting, the NCA achieves a mean accuracy of 72.17 ± 35.68, compared with 67.96 ± 44.11 for the WL baseline. Thus, the two methods are broadly comparable in classification accuracy and robustness, with the NCA showing slightly stronger aggregate robustness in the damaged-cell setting.

Their main difference lies in distributed execution. The NCA reaches its prediction through a fixed number of purely local update cycles, whereas the distributed WL

baseline incurs additional communication overhead that grows with the diameter of the shape and, in its exact form, requires repeated global coordination (Fig. S6). Moreover, the WL baseline is fundamentally a graph-feature extraction and classification method, and therefore does not directly provide the local directional signals required for damage localization and regrowth. Thus, the main advantage of the NCA is not only that it achieves strong classification and robustness, but that it is, to our knowledge, the only method that does so in a form that is naturally compatible with fast decentralized execution and can be extended to fault localization and recovery.

3) The Conclusion could be strengthened in two ways. First, the authors should specify potential application scenarios for this technology. Second, a discussion of the gaps between the system and real biology (e.g., in energy acquisition, cellular differentiation, or scale invariance) would better outline future research challenges and directions.

In the revised Discussion, we now make the application relevance more concrete across different time horizons. In the short term, we highlight interactive edutainment and "smart LEGO"-like systems, in which users assemble structures from intelligent blocks that can collectively infer the emerging shape, provide local visual feedback, detect assembly mistakes, and indicate where additional blocks should be placed. For example, a child could assemble a rough animal shape from smart blocks, and once the collective recognizes it as, for instance, a dog, duck, or cow, the system could respond with the corresponding animal sound or other interactive behaviors.

In the near term, we point to smart modular surfaces and architectural components—such as facade tiles, wall panels, or other reconfigurable assemblies—where local communication could support distributed fault detection, structural self-assessment, and human-guided maintenance without centralized monitoring. In the longer term, as the hardware matures through improved actuation, energy handling, specialization, miniaturization, and connector design, the same decentralized framework could become relevant to adaptive smart materials, damage-tolerant modular robotic structures, autonomous construction, remote inspection, and space systems.

At the same time, we now state more clearly how far the present system still is from biological systems, including the lack of autonomous energy regulation, richer long-term specialization, and the sensing, actuation, and material adaptation needed for continuous growth and repair. We also connect these gaps directly to future work on closed-loop growth, stronger recovery, and scaling from current prototypes and simulation to larger physical systems.

4) The results section has an inconsistent narrative flow. The authors describe Figures 2A-D, jump to Figure 3, and then return to Figures 2E-H. We recommend presenting all of Figure 2 (A-H) together to first fully describe the classification capabilities, followed by the robustness analysis in Figure 3. This will significantly improve the logical coherence.

We thank the reviewer for this useful suggestion, which indeed made the flow of the paper much better. We now first show all the classifications results and then the robustness analysis.

Reviewer #2 (Remarks to the Author):

This work presents the first 3D hardware realization of decentralized morphogenesis and self-recognition. Its most significant contribution is the physical demonstration of a collective of nearly 200 independent modules achieving consensus on their 3D global shape (e.g., “plane,” “table,” “boat”) without centralized control or global positioning, relying solely on local communication. The 3D convolutional design, featuring a cross-shaped kernel, is well-motivated and draws inspiration from biological systems. The use of PCBs as structural components is both elegant and cost-effective.

However, the experimental setup and evaluation of results suffer from notable limitations. While the physical hardware implementation is constrained to a lower resolution (15×15×15) due to practical considerations, all experiments on robustness and damage recovery were conducted entirely in simulation. From a computer graphics research perspective, such experiments could—and arguably should—be performed at higher resolutions (e.g., 64×64×64 voxels) to yield more meaningful and generalizable insights. The current results do not provide sufficient evidence that the NCA-based approach will perform similarly at higher resolutions, limiting the scalability claims.

We thank the reviewer for this important comment. We note that we addressed a closely related concern in our response to Reviewer 1, but we repeat the main points here for completeness.

We agree that scalability cannot be supported convincingly from the original low-resolution setting alone. In response, we expanded the manuscript with additional experiments at substantially higher resolutions, evaluating the same framework on both 32 × 32 × 32 and 64 × 64 × 64 grids. These new experiments go well beyond the hardware-constrained regime and include much larger and more intricate voxelized structures, including shapes with hollowness, internal cavities, spiral-like structure, and assemblies of more than 18,000 cubes.

The new results show that the NCA-based approach remains effective at higher resolutions, especially when model capacity is increased. With 128 hidden channels, classification accuracy remains high across all three resolutions, reaching 0.98, 0.96, and 0.96 for the 15^3, 32^3, and 64^3 settings, respectively. We also now report training efficiency across resolutions: the models exceed 90% accuracy after 108, 517, and 611 training steps for the 15^3, 32^3, and 64^3 settings, respectively. These results provide direct evidence that the framework is not restricted to the original low-resolution setting and

instead exhibits encouraging scaling behavior as resolution increases, provided sufficient model capacity.

We have revised the manuscript accordingly and now present the higher-resolution experiments explicitly as evidence that the decentralized NCA framework extends to larger, denser, and structurally richer voxelized objects.

Furthermore, the self-recognition and damage detection algorithms are implemented as separate processes, despite both being based on the same NCA framework. Notably, damage detection requires conditioning the NCA on a given shape class. If biological systems can simultaneously exhibit both capabilities, it is reasonable to expect that the neural cells in this system should also be capable of concurrent self-recognition and damage detection. Since the experiments already demonstrate robust classification under 5–15% fault conditions, extending the framework to detect and localize these faults within the same process would be a more significant and coherent step toward true decentralized self-organization.

We agree that, in a biologically inspired decentralized system, it is important to examine whether self-recognition and damage detection can occur concurrently within the same NCA process. In the revised manuscript, we now explicitly include such a model: a self-classifying NCA model that simultaneously predicts shape class and local damage direction for each cell. This addresses the reviewer's concern directly by showing that concurrent self-recognition and damage detection are indeed possible within a single decentralized learned dynamics.

At the same time, we found that the joint and shape-conditioned formulations are complementary rather than redundant. The joint model is the more biologically natural formulation, because it allows cells to infer both global identity and local damage from morphology alone. However, this same property creates a limitation under severe damage: when the shape changes enough to become morphologically ambiguous, the class prediction itself can shift, which in turn can hinder recovery (see Fig. 5C). We now discuss this explicitly in the manuscript, for example in cases where a heavily damaged chair can become more consistent with a table-like morphology.

For this reason, we also retain the shape-conditioned model as a complementary extension (Figure 5). Its role is not to replace concurrent self-recognition and damage detection, but to provide a recovery mechanism for regimes in which morphology alone is no longer sufficient to infer the intended structure uniquely.

Together, these additions strengthen the manuscript in two ways: they show that simultaneous self-recognition and damage detection can emerge within a single NCA framework, and they clarify why an additional shape-conditioned variant is still useful for more difficult recovery settings.

Reproducibility is also hindered by insufficient detail in several key areas:

1. Cell State Initialization at t=0: The manuscript does not specify the initial values for the 28 state channels (hidden, logits, alpha), which are critical for reproducibility.

We now clarify that the state channels are initialised with all zeros.

2. Inter-cell Communication Content: The manuscript lacks a clear description of what "information" is aggregated from neighboring cells. Do cells exchange the full 28D state vector, a subset of it, or some processed perceptual features?

We agree that the communicated content between neighboring modules was not described clearly enough. We have now clarified in both the main text and the Methods that modules exchange their current full NCA cell state, not intermediate perceptual features. Specifically, each message contains the full 28-dimensional state vector (1 alpha channel, 20 hidden channels, and 7 class-logit channels), which is sent to connected neighbors and used as input to the next local update.

3. Consistency in Damage Identification: The algorithm does not explain how consistency is enforced among different cells during damage identification. For example, if two normal cells, A and B, are both adjacent to a region C, and cell A identifies C as damaged while cell B does not, what mechanism resolves this conflict?

We have now clarified that damage identification is not taken from a single neighboring cell in isolation. Instead, for each cell, we first sample candidate damage-direction labels from the multinomial distributions defined by the direction logits received from its neighbors. The cell then assigns its final damage classification by majority vote over these neighboring sampled outcomes. Thus, if adjacent cells disagree about whether a nearby region is damaged, the conflict is resolved locally through neighborhood voting rather than by relying on a single cell's prediction.

These omissions significantly impede faithful re-implementation and validation of the proposed system.

In summary, while the physical demonstration of decentralized shape classification is impressive, the work remains preliminary. The reliance on low-resolution simulations, the separation of recognition and fault detection, and the lack of implementation details prevent a full assessment of scalability and robustness. Given these unresolved core issues, I recommend a major revision. Essential revisions should include:
1. Experiments conducted at higher resolutions;
2. A unified algorithm enabling simultaneous shape recognition and fault detection;
3. Clear and complete specifications for initialization, inter-cell communication, and conflict resolution mechanisms.

We believe we have now addressed all these important points brought about by the reviewer in our revised manuscript, which significantly increased the experimental setup and its evaluation.

Reviewer #3 (Remarks to the Author):

Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery

This paper introduces the concept of a 'smart brick', which contains its own local sensing, compute – predominantly used to run a local neural network - and communication, making each brick an autonomous agent in its own right. The bricks can be combined into a wide variety of morphologies.

The authors run a good variety and depth of experimentation, characterising emergent system properties including shape classification, fault tolerance, and shape repair. Compared to the state of the art, these bricks implement full decentralisation, with experimentation demonstrating capabilities that exceed state of the art.

Presentation is pretty good, it feels like a Transactions on Graphics paper. Figures are useful, well-captioned, and clearly communicate key features of the proposed approach and related experimentation.

Thank you!

The authors should comment on physical scalability – they simulate hundreds of blocks, and systems like this make most sense with a large number of blocks, e.g. into the hundreds. The blocks themselves are made out of PCBs and seem to require manual assembly e.g., to cable tie the 'lid' on. How long does it take to create one block, and how much manual effort is required? How expensive is a block? It seems like fabrication may be a bottleneck to any practical realisation.

We thank the reviewer for this thoughtful and important comment. We agree that physical scalability is a central issue for systems of this kind, and we have revised the manuscript to address it more explicitly and quantitatively.

First, in the revised version, we now report the practical fabrication details of the current hardware platform. A brick is built using assembled PCB. The main manual effort comes when arranging the faces and soldering the joining contacts together. Currently, with no more optimization than a jig to hold the loose faces in place, the process takes around 10 minutes per brick (including connecting the internal cable and placing the brick cover around) mainly due to the soldering time (see movie S5).

Importantly, this process can be optimized, for example by using a custom soldering tool or an automatic solder feeder, which would bring the assembly time down to around one minute. Regarding the price, given current offers by the PCB manufacturer and ordering parts for 100 bricks, one brick costs around 20 EUR to produce. Ordering parts for 1,000 bricks brings the price down to 15 EUR per brick. These costs are inflated by our use of off-the-shelf microcontroller and DC-DC converter modules, which were chosen to reduce development time for this proof-of-concept platform rather than to optimize for large-scale manufacturing.

A more scalable hardware revision would replace these modules with the corresponding integrated circuits (ICs) and supporting discrete components directly on the main PCB. Doing so would reduce component cost, decrease part count and assembly complexity, improve packaging density, and make the design more suitable for automated manufacturing at the scale of thousands of units. The trade-off is that such a redesign would require substantially greater engineering effort, including custom PCB development, power-management design, RF/layout optimization, and additional testing.

To address the reviewer's broader concern about whether the approach itself remains meaningful at larger scales, we have also strengthened the manuscript with new higher-resolution simulation results. Beyond the hardware experiments with 200 bricks, we now show that the same framework extends to 32 × 32 × 32 and 64 × 64 × 64 settings, including much larger and more intricate structures with up to roughly 18,000 cubes. Taken together, these results suggest that the current limitations are less a fundamental barrier of the decentralized approach than an opportunity for future hardware development. In other words, the algorithmic framework already appears compatible with substantially larger collectives, and advances in fabrication, integration, power handling, and automated assembly could enable these capabilities to be realized in much larger physical systems.

A big issue here is the practical application. The motivation given seems to be predominantly around 'modular robotics' and 'cellular repair', but there is no actual use case provided. Are the blocks to make up a robot? Are they to repair a larger structure? What structure could benefit from decentralised self-repair and self-classification, and be essentially made of PCBs at the given scale? Practical applicaiton seems to be well out of reach of the current technology, which should be highlighted more forthrightly in the conclusions if I am correct.

We agree that the original manuscript did not make the practical application sufficiently concrete, and that it also did not clearly distinguish between what is demonstrated by the current hardware and what remains a longer-term vision. We revised the paper to make this distinction more explicit.

In the revised framing, we see the practical relevance of the work across three timescales. In the realistic short term, the most immediate use case is interactive edutainment (i.e. "smart Lego"), where users assemble structures from intelligent modular blocks and the collective infers the emerging shape through local

communication alone. In this setting, the system could provide light or sound feedback, identify what has been built, and guide the next assembly step by indicating where additional blocks should be attached (e.g. depending on the type of animal a child has built, it makes a different sound). This application is consistent with the current hardware capabilities. [In fact, we have recently been in talks with Lego about a collaboration and they just introduced a smart bricks system that uses a centralized approach (https://www.lego.com/en-us/smart-play/article/innovation).]

In the near term, we see opportunities in smart modular surfaces, such as facade tiles, wall panels, or other distributed assemblies that can locally detect missing or damaged elements and support human-guided maintenance. In these scenarios, the same decentralized shape inference and damage-localization mechanisms demonstrated here would function as distributed monitoring and repair-guidance tools.

In the far term, the broader vision is adaptive smart materials and self-reconfigurable modular robotic systems capable of autonomous repair and morphological adaptation.

At the same time, we agree with the reviewer that the latter vision is not yet achieved by the present platform. The current system is best understood as an important physical proof of concept for decentralized morphology inference and damage localization, rather than an application-ready self-repairing material or robot. As described in the manuscript, the modules are centimeter-scale PCB-based cubes, externally powered, and do not yet include the actuation, autonomous energy handling, or material integration required for active self-repair or structural deployment. The Discussion now notes these limitations and points to actuation, miniaturization, improved connectors, and closed-loop growth as important future directions.

I am also interested in how small such bricks could be made – do the authors have any comments on this?

The present modules measure 44 × 44 × 47 mm and were intentionally designed to prioritize rapid development, robustness, and experimental validation over miniaturization. As mentioned above, the current implementation relies on pre-assembled wireless microcontroller and DC-DC converter modules, identical side PCBs, internal cabling, and a mechanically simple enclosure strategy. These choices reduced development time and risk, but they also introduce substantial overhead in size and packaging.

We believe that a substantial reduction in module size is feasible in a future hardware revision. A conservative redesign, aimed at preserving roughly the same functionality and full 3D six-face connectivity, could likely reduce the cube dimensions to approximately 20–25 mm. This level of reduction could be achieved by replacing the current wireless and power modules with the corresponding integrated circuits and discrete components directly on the PCB, eliminating the internal cable in favor of

board-to-board interconnects, redistributing components more efficiently across the cube faces, and redesigning the LED and routing layout. Importantly, this would still correspond to a system with comparable sensing, communication, and computation capabilities to the present one, but in a much denser form factor.

An even more aggressive custom redesign could likely push the size further, potentially into the 15–20 mm range. However, we view this as requiring significantly greater engineering effort and a more specialized hardware-development process. At that scale, the main challenges would no longer be just component selection, but also RF layout for the wireless subsystem, tighter power-delivery design, more demanding inter-board interconnect strategies, and significantly more constrained packaging and assembly tolerances. In other words, miniaturization at that level appears plausible, but it would require moving from an accessible proof-of-concept prototype to a highly optimized custom electromechanical design.

Actuation is a major issue. I recognise that the authors discuss this in the 'limitations' section, which is good. However, most of the capabilities demonstrated in the simulation cannot currently be transferred to the real agents. This compounds my above issue with practical application and fabrication scalability and suggests that the physical block systems of reasonable size and complexity will be prohibitively difficult to create.

We agree that actuation is likely the component that will require the greatest additional hardware development. While our current NCA-based system already supports shape classification and damage detection in hardware, fully automated damage recovery presents a substantially greater challenge.

In the longer term, the required actuation layer could build on ideas demonstrated in self-reconfigurable robotics. For example, the modular robotic system Roombots [1] achieves reconfiguration through active latching and rotational degrees of freedom, while ATRON [2] shows how lattice-based modules with local communication, mechanical connectors, and simple revolute joints can self-reconfigure in three dimensions. Of course, these systems are themselves primarily demonstrations of concept rather than fully mature, scalable platforms. Nevertheless, they illustrate that the core mechanisms for physical reconfiguration already exist, suggesting that the main limitation is not the decentralized inference framework introduced here, but its integration with sufficiently advanced and scalable actuation hardware.

[1] Spröwitz, Alexander, et al. "Roombots: A hardware perspective on 3D self-reconfiguration and locomotion with a homogeneous modular robot." ***Robotics and Autonomous Systems*** **62.7 (2014): 1016-1033.**

[2] Jorgensen, Morten Winkler, Esben Hallundbk Ostergaard, and Henrik Hautop Lund. "Modular ATRON: Modules for a self-reconfigurable robot." ***2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566).*** **Vol. 2. Ieee, 2004.**

It would also be helpful to know how sturdy and robust these bricks are – can the authors perform some testing with significant numbers of cycles of assembly/disassembly, even for two bricks being repeatedly connected and disconnected, to get some idea of robustness and implications for reusability?

Repeatedly connecting and disconnecting the bricks over more than 100 cycles did not result in any observable connection issues, consistent with the manufacturer's specifications, which rate the pin headers and connectors for at least 100 mating cycles without loss of electrical performance. While the manufacturer does not specify an upper limit, we did not experience connector-related failures or degradation in any of our experiments, which involved numerous cycles of assembling and disassembling modules, building, and disassembling shapes. Even more robust alternatives such as spring-pin–based connector systems used in other modular robotics have been shown to support thousands of cycles without significant wear or damage [4].

The core limitation with the damage detection is that it only works on a specific instance of a given class, rather than across a class. It would be good to see what happens when this is done per class, and what the failure modes are – e.g., what makes a table repair itself into a chair, and can chains of local activations be somehow visualised? This would be a very interesting study to see, as you could track the emergent behaviour over time; e.g., does a broken table start to repair itself as a table, then switch to a chair? Under what conditions does it switch between classes? It must also be said that this combines with the rest of the limitations in this study to somewhat reduce the impact of the paper overall.

We now address this core limitation with a new approach, which we call self-classifying recovery; this method performs class recognition and damage detection simultaneously within the same decentralized dynamics (see results in Figure 5). The self-classifying recovery approach directly addresses the reviewer's concern, because damage detection is no longer tied only to a specific target instance: each cell must infer both the class identity and the local damage direction from morphology alone. In this sense, the revised manuscript now explicitly studies damage detection and recovery at the class level, rather than only in a shape-conditioned setting.

Second, we now clarify the failure mode that motivated retaining the original shape-conditioned variant as a complementary model. In the joint model, recovery is coupled to the current class prediction. As a result, if damage removes the structural features that distinguish one class from another, the inferred class can change during the recovery process. This is precisely the kind of behavior highlighted by the reviewer: for example, a severely damaged chair can become more consistent with a table-like morphology, and the collective can then switch from repairing it as a chair to treating it as a table (Figure 5C). We now discuss this class-switching behavior explicitly in the revised manuscript and use it to motivate the shape-conditioned model, whose purpose is to preserve recovery toward a particular intended shape when morphology alone becomes ambiguous. Both approaches show high performance for their particular use case.

More broadly, we agree that these emergent transitions are scientifically interesting rather than merely failure cases. The revised manuscript therefore treats the two models as complementary. The self-organized joint model is the more biologically natural formulation, because it allows self-recognition and damage detection to arise together from local interactions. The shape-conditioned model is introduced to resolve ambiguity in regimes where severe damage makes class identity underdetermined. Together, these additions allow us to address both the class-level setting requested by the reviewer and the associated failure modes, while clarifying the trade-off between flexible self-organization and target-specific recovery.

Despite this, there are some solid contributions. Self-repair and fully distributed classification (which even somewhat holds across different scales and is therefore scalable!) are fundamental features of any successful future modular robotic/structural system. The authors draw on some biological inspiration, but it is unclear to me if biological self-repair works anything like the mechanisms proposed in the paper. If this biological inspiration has merit, it should be expanded on and concretely demonstrated. If not, it should be rewritten to appropriately represent the extent to which the proposed approach is bio-inspired.

We thank the reviewer for this thoughtful comment. We agree that the original manuscript could have more clearly distinguished between biological inspiration at the level of high-level design principles and any claim of mechanistic similarity to biological self-repair. Our intention is the former, not the latter. Accordingly, we revised the manuscript to better represent the extent to which the approach is bio-inspired.

Specifically, we rewrote several passages in the manuscript to remove wording that could imply a direct biological blueprint or mechanistic correspondence (e.g., terms such as “blueprint,” “mirror,” and similar formulations). Throughout, we now state more explicitly that the proposed system is not intended as a mechanistic model of biological regeneration, but is instead inspired by the broader principle that robust global behavior can emerge from decentralized local interactions.

We also strengthened the Discussion to make this distinction clearer by explicitly highlighting both the biological motivation and the remaining gaps between the present system and living systems. In particular, we now emphasize that, unlike biological tissues, our platform does not yet include autonomous energy handling, richer specialization, sensing, actuation, or material adaptation for active repair. The contribution of the paper is therefore not a biological model of self-repair, but a physical demonstration that decentralized shape classification and damage-aware behavior can be achieved in a modular system through local communication and identical distributed control.

We appreciate the reviewer’s positive assessment of the distributed classification and self-repair contributions, including the observation that the approach shows some robustness across scale. In the revised manuscript, we aimed to preserve that core contribution while presenting the biological framing more carefully and accurately.

Minor typo “no more damaged was detected.”

Fixed.

Overall I’m torn on this one. It’s certainly interesting, well written, and well presented. It has significant limitations that lessen impact and, in particular, stop the core potential contributions being practically realised. Solving these limitations would involve significant effort in hardware development but would improve the article and make it a very high quality Nature Comms article. I'm hesitant to advocate for acceptance without the above-mentioned additions being strongly considered by the authors.

We thank the reviewer for this balanced assessment. We agree that important limitations remain, particularly in actuation and large-scale physical self-reconfiguration. At the same time, we view the present work as an important stepping stone: it establishes, in physical 3D hardware, decentralized shape classification and damage-aware behavior as viable ingredients for future modular robotic and adaptive material systems. In the revised manuscript, we therefore aimed to strengthen the paper along the reviewer’s suggested directions (through additional scalability experiments, a joint self-recognition and damage-detection model, a decentralized baseline comparison, and a clearer discussion of practical scope and current limitations), while being explicit that fully realized self-reconfiguring systems remain a longer-term goal. We hope these revisions make the manuscript’s contribution and its role as an important foundation for future hardware advances clearer.

Author response "Smart Cellular Bricks for Decentralized Shape Classification and Damage Recovery"

Reviewer #1 (Remarks to the Author):

The authors have addressed my comments

We are glad we could address all of your comments.

Reviewer #2 (Remarks to the Author):

The authors have adequately addressed the concerns raised in the previous round. The revised manuscript now includes higher-resolution experiments (32^3 and 64^3), a joint self-recognition and damage detection model, a WL baseline comparison, and implementation details. These additions satisfactorily resolve my earlier concerns about scalability, concurrency, and reproducibility.

We are happy to hear we could address your previous comments.

The manuscript in its current form presents a clearer and more complete case for its contributions. I have only a few minor suggestions regarding presentation:

1. Narrative flow: The WL baseline comparison and high-resolution results currently appear in the Discussion. I recommend moving them into the Results section, so the narrative flows more smoothly.

We now moved them to the result section.

2. Figure 2H: On page 17, the text refers to “Figure 2H,” but subpanel H within Figure 2 appears to be missing or has not been updated. Please verify.

Fixed.

3. Consistency in supplementary figures: In Figures S5–S7, the color for the same algorithm varies across panels. A consistent color assignment for each method would make comparisons easier.

We now use consistent color assignments for these figures.

Overall, the revised work is substantially stronger. The remaining issues are minor and pertain only to presentation. Once addressed, I believe the manuscript is suitable for publication.

Thank you!

Reviewer #2 (Remarks on code availability):

The .npy and .zip data files at the root level would be better placed in a dedicated subfolder, and adding a .gitignore to exclude intermediate files such as .vscode/ and __pycache__/ would keep the repository cleaner.

Fixed.

Reviewer #3 (Remarks to the Author):

The manuscript has been improved and many of the limitations have been either addressed or more specifically called out. The justification vs. biology is now fairer.

Thank you for your previous comments that helped to significantly improve our manuscript.

I appreciate the scaling baseline, however the computational cost of scaling in simulation must be explicitly called out. This needs adding to the manuscript.

Otherwise I am happy that the paper is now improved. Given the editor is happy with the fundamental limitations of the approach and questionable real-world applicability that I previously mentioned, this paper - once computational cost is added - is good enough to be accepted.

We have now added more details on the computational cost in the discussion section.