Architectural swarms for responsive façades and creative expression

Merihan Alhafnawi[1*], Jad Bendarkawi[2†], Yenet Tafesse[3†],
Lucia Stein-Montalvo[4], Azariah Jones[5], Vicky Chow[6],
Sigrid Adriaenssens[7], Radhika Nagpal[1,3]

[1]Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[2]Department of Electrical and Computer Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[3]Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA.

[4]Department of Civil and Environmental Engineering, Northwestern University, Evanston, Illinois 60208, USA.

[5]Department of African American Studies, Princeton University, Princeton, New Jersey 08544, USA.

[6]School of Architecture, Princeton University, Princeton, New Jersey 08544, USA.

[7]Department of Civil and Environmental Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[*]Corresponding author. Email: m.alhafnawi@princeton.edu

[†]These authors contributed equally to this work.

Living architectures, such as beehives and ant bridges, adapt continuously to their environments through self-organization of swarming agents. In contrast, most human-made architecture remains static, unable to respond to changing climates or occupant needs. Despite advances in biomimicry within architecture, architectural systems still lack the self-organizing dynamics found in natural swarms. In this work, we introduce the concept of architectural swarms; systems that integrate swarm intelligence and robotics into modular architectural façades to enable responsiveness to environmental conditions and human preferences.

We present the Swarm Garden, a proof-of-concept composed of robotic modules called SGbots. Each SGbot features a buckling-sheet actuation, sensing, computation, and wireless communication. SGbots can be networked into reconfigurable spatial systems that exhibit collective behavior, forming a testbed for exploring architectural swarm applications. We demonstrate two application case studies. The first explores adaptive shading using self-organization, where SGbots respond to sunlight using a swarm controller based on opinion dynamics. In a 16-SGbot deployment on an office window, the system adapted effectively to sunlight, showing robustness to sensor failures and different climates. Simulations demonstrated scalability and tunability in larger spaces. The second study explores creative expression in interior design, with 36 SGbots responding to human interaction during a public exhibition, including a live dance performance mediated by a wearable device. Results show the system was engaging and visually compelling, with 96% positive attendee sentiments. The Swarm Garden exemplifies how architectural swarms can transform the built environment, enabling "living-like" architecture for functional and creative applications.

Architectural swarms fuse swarm robotics and architectural façades to create spaces that adapt to people and the environment.

INTRODUCTION

Living architectures, such as plants, beehives, and army ant bridges, are constantly evolving in response to their environments with immediate and long-term adaptations. These adaptations emerge from self-organization, whereby individual agents (for example, cells and insects) interact locally with each other and sense their local environment, and through networked interactions, create globally complex responses (*1*). Plants optimize their shape to sunlight and nutrients through mechanical and signaling interactions of individual cells that sense external stimuli (*2*). Beehives continually adapt the distribution of brood vs. honey cells through self-organized actions of individual bees (*1*). Army ants build living bridges out of their own interlinked bodies that adapt in real-time to traffic and terrain changes (*3*). These natural, living architectural systems have inspired the field of swarm

intelligence, which explores how natural agents self-organize to create complex collective behaviors and how such algorithms can be applied to the design of artificial systems (*4*).

In contrast to these living architectures, human-designed architectures are for the most part static, designed with rigid elements that remain fixed even as the environment and occupants' needs change (*5*). This inflexibility hinders the ability to adapt to different occupants' requirements or daily, seasonal and yearly climate variations. However, if human architectures were designed to be dynamic, or "living-like", they could incorporate many interacting agents that self-organize to respond to daily-to-seasonal climate variations, different occupant requirements, and even enhance health and well-being. Such agents could leverage swarm intelligence to create spaces that evolve over time.

Currently, the application of swarm intelligence in architecture has been limited to the design process to address "task fitness" challenges. The architectural design process is complex, involving multi-objective optimization. Its task fitness refers to how well a physical design balances often competing requirements such as layout, aesthetics, structural and environmental performance, and material efficiency. Traditional architectural design processes rely heavily on tacit knowledge and human expertise, limiting the exploration of diverse solutions. Swarm intelligence algorithms have been successfully applied to a multitude of complex optimization scenarios (*6, 7*), enabling the exploration of large complex design solution spaces and providing solutions that are not apparent in the traditional design process. For example, swarm algorithms inspired by plants, insects, and bird flocks have been applied to creative structural design and complex energy optimization scenarios (*6, 8, 9*). However, unlike the swarm systems that inspire them, these architectural systems are typically static.

Adaptive and dynamic architecture has been explored in the design of façades. The work exploring biohybrid techniques in architecture (*10*) aims to create "living buildings" that integrate biological elements. Examples include the integration of algae panels for shading and energy generation (*11*), and the use of robotic modules to guide plants' growth to create a structure (*12*). These systems are limited to environments that can sustain living materials and do not integrate human interaction. However, they inspire the design of façades that mimic natural beauty while also being functional. Most adaptive façade systems typically consist of arrays of mechanical modules that are responsive to external stimuli in a physical environment (*13*). The overarching goal is

climate-adaptive building skins that respond to changing conditions or occupant preference – often by shape-morphing – to improve thermal or visual comfort, ventilation, or acoustic properties (*14, 15, 16, 17, 18*). Modular designs allow for tunability in mixed-use spaces. The mechanical modules are typically based on rigid mechanics of origami or rigid-body motions (*19, 20, 21, 22*), and famous examples include Al Bahr towers (Aedas Architects, Abu Dhabi, United Arab Emirates) and Institut du Monde Arabe (Jean Nouvel, Paris, France). These modular façades are often limited by the mechanical complexity and lack of robustness of the individual rigid modules. This has led to the recent exploration of soft modules with continuous deformation, which bring their own challenges in materials and modeling (*13, 23*). Notable examples include façades comprised of 4D-printed hygromorphic bilayer panels (*24*), thermally responsive bimetallic beams and shells for temperature and ventilation control (*5, 25*), and torsional-compressive buckling driven lamellar shading elements such as the One Ocean pavilion (soma architecture, Yeosu-si, South Korea). These soft façades are often inspired by nature - for example, the shape-change and phototropic motions of plants - and incorporate biomimicry. However, none of these adaptive architectural systems incorporate swarm intelligence. Although modular, they are often rigid, with control that is either completely centralized (hard to scale) or completely decentralized (risking uncoordinated or conflicting actions). They often do not allow programmability to adapt to new environments or to incorporate human interaction. In contrast, swarm robotic modules can deliver globally coherent behavior and enhance robustness against individual errors through redundancy. They also provide programmability and modularity, enabling reuse and adaptation to new tasks, integrating human preferences through human-robot interaction, and deployment in a wide variety of environments. This reusability supports sustainability, which further could be enhanced by reduced costs from mass manufacturing and renewable energy use.

The field of swarm robotics applies concepts from swarm intelligence to the design of physical systems, where many robotic agents cooperate to exhibit complex behaviors that are scalable and robust (*4*). Traditionally, most swarm robots have been developed as mobile ground robots, aerial robots such as quadcopters, or underwater robots. Example applications include inspection, agriculture, and logistics (*26*). In architecture, robots have been applied primarily to the construction process (*27, 28, 29, 30*), although only a few leverage swarm intelligence (*30*). There has also been an exploration of self-reconfiguring modular robots for furniture and interior design (*31, 32*).

Swarms specifically designed to interact with humans are uncommon, and usually focus on a single application such as education (*33, 34*) or art (*35, 36, 37*). Notable exceptions include systems like MOSAIX (*38*) and Zooids (*39*); MOSAIX uses mobile ground robots to help human groups collaborate and Zooids uses table-top mobile robots as a display for social tasks. These systems demonstrate the potential of swarm robotics to move beyond traditional applications and engage directly with humans in diverse, meaningful ways. However, their form is still restricted to wheeled mobile robots. Swarm robots that are integrated into human environments - such as architectural façades – are an emerging physical design space that remains to be explored and represents an exciting direction for swarm and human-robot interaction research.

In this article, we introduce the concept of architectural swarms—systems that integrate swarm intelligence and swarm robotics into modular architectural façades to achieve complex environmental and human-centric adaptation through self-organization. Such adaptations could include improving ideal natural light levels and distribution in human spaces, responding to diverse user comfort preferences, and creating visually appealing interior environments (*40, 41*); all factors of which could contribute to uplifting mental moods and increasing overall occupant well-being (*42, 43*). It could also enable new forms of creative expression. Architecture has served as a medium of human creativity and expression, with designers weaving aesthetics and culture alongside functionality to communicate with the world (*44, 45*). In a similar way, robotics has been explored as a medium for creative expression (*46*). Robot swarms have been used to to paint (*35*), play music (*47*), and dance (*48*), helping humans convey emotions and ideas beyond purely industrial applications. To this end, we created the Swarm Garden—a proof of concept of architectural swarms—that enables exploration of adaptive shading and creative expression through a reconfigurable swarm of robot modules, called SGbots, which collectively respond to light, motion, and neighboring robots (see Fig. 1A–D and movie 1). The system can be configured into large spatial networks that coordinate through programmable nearest-neighbor rules. Each SGbot functions as a novel façade module whose soft, continuous buckling-based blooming-like actuation achieves complex three-dimensional shape change using only a single degree-of-freedom actuator. It offers organic motion with minimal mechanical complexity, replacing complex folded or hinged origami structures that are difficult to manufacture and prone to jamming failures. In our first case study, we show that architectural swarms can perform self-organizing adaptive shading, where an opinion-

dynamics controller robustly enables collective adjustment to sunlight in a 16-SGbot real-world window deployment (Fig. 1E), with additional tunability demonstrated in large-scale simulations. In our second case study, we demonstrate the potential of architectural swarms as expressive and interactive media for interior design, supported by a public exhibition where more than 100 participants and a dancer (Fig. 1F) engaged with the Swarm Garden through various interaction modes, including a wearable device. Through these case studies, the Swarm Garden guides the deployment of "living-like" architectures for functional and aesthetic applications, demonstrating the versatile potential of architectural swarms. It advances adaptive façade design by incorporating swarm intelligence and integrating multiple inputs for dynamic responses, and advances swarm robotics by pushing the boundaries of human-centric applications.

RESULTS

This section presents the design of the Swarm Garden and results from real-world deployments for two case studies. Case Study 1 investigated the application of the swarm as an adaptive shading façade. Case Study 2 investigated the application of the swarm for interior design and creative expression.

The Swarm Garden system design

The Swarm Garden consisted of 40 modular, re-arrangeable, robotic agents called SGbots. These robots formed a responsive architectural swarm through an interconnected WiFi network, enabling a shared communication protocol that facilitated collective decision-making among the robots and interaction with other devices in their environment, such as sensors and wearable devices. The Swarm Garden could be configured and deployed in different spatial configurations, as shown in Fig. 1E-J. A localization method, detailed in Materials and Methods, allowed the SGbots to identify their nearest neighbors, enabling swarm intelligence algorithms and emergent collective behavior.

Each SGbot was an autonomous robot, capable of sensing, computation, actuation and communication. It had a back-facing ambient light sensor that allowed it to sense the intensity of light, and a front-facing proximity sensor (the "interaction sensor") that allowed it to respond to human interaction (see Fig. 2). The response of each SGbot after sensing the environment and communicating

was through “blooming”. This process involved pulling a centrally-clamped thin plastic sheet, the “plate”, through a circular opening, actuated by a simple single degree-of-freedom mechanism using a threaded rod that retracted and extended the plate. We refer to the process of the sheet buckling (when it is retracted) and then flattening (when it is extended) as “blooming”. This mechanism was chosen for its high simplicity, reliability, and precision. Other forms of actuation were considered, and we elaborate on their trade-offs in the Supplementary Methods. A time-of-flight sensor was positioned near the nut of the threaded rod, acting as an encoder to measure and monitor how far the plate was extended, denoted by “encoding sensor” in Fig. 2A. The plate confinement caused the plate to buckle circumferentially into flower-like patterns of truncated developable cones (see Figs. 2B,C), as first examined by Stein-Montalvo et al. (*49*) and further rationalized by Seffen (*50*). The Swarm Garden leveraged the controllable buckling patterns and large deformations exhibited by the confined sheets to enable modular adaptive façade applications. The blooming deformation had a “gentle” appearance, and took 10 seconds to go from fully flattened to fully buckled. Faster actuation speeds were possible, but they required more power - highlighting a trade-off between speed and energy efficiency. To further enable fast dynamic appearance change, Light-emitting diode (LED) strips were mounted on the back-side acrylic sheet of the robot. This design feature allowed the robot to change colors (see Fig. 2D), offering ease and flexibility to present multiple looks, especially for artistic and aesthetic purposes. See Fig. 2E for the process of blooming, Materials and Methods for more details on the design of the SGbot and table S1 for each SGbot’s bill of materials.

The Swarm Garden was highly versatile and robust; it was tested in multiple deployments including an office environment, a public exhibition in the Lewis Center for the Arts, Princeton University, with over 100 attendees, as well as robotics conferences in New York City and Detroit. Many collective behaviors were successfully implemented, including classical swarm intelligence algorithms such as gradients and opinion dynamics.

Case study 1: Swarm robots for adaptive façades

The Swarm Garden could be envisioned as a light-adaptive shading façade, where the individual agents sense external light in their spatial region and collaborate to create user-tailored internal

illuminances (Fig. 1H-J). In this case study, we aimed to understand how self-organization can enable a façade's capacity to respond to changing external lighting and weather conditions, mitigate sensor and communication failures, and integrate user preferences. We studied this with the physical deployment of 16 robots in an office window over an 8-day period, where there was substantial day-to-day weather variation. We enhanced individualistic response by a collective decision-making approach that attempted to maintain a desired internal illuminance and adapt to failures. We also investigated a more complex scenario - a large-scale atrium ceiling swarm, using numerical simulations (Rhino3D/Grasshopper). Additional analyses of both the collective decision-making and atrium studies are provided in the Supplementary Methods.

Adaptive shading with an individualistic approach

For three days, 16 SGbots were placed on an office window (see Fig. 1E), operating continuously for seven hours each day, from 7:00 to 14:00. The experiments were conducted in mid to late August, a time when the sun typically appeared at the chosen window around 7:15 a.m., intensified by approximately 9:00 a.m., and disappeared from the window's view by about 11:00 a.m (obscured by a nearby building). To assess the systems' responsiveness to sunlight, we developed a proportional controller that adjusted the SGbot's bloom level (the plate displacement) to the illuminance readings from the ambient light sensor. The aim of the proportional controller was to cause the plate to fully extend, blocking sunlight when it was strong. As the sunlight weakened and the room darkened, the plate would gradually buckle to allow more light to pass through. The plate displacement was measured by the encoding sensor placed on the nut adaptor of the threaded rod (see Fig. 2A). The ambient light sensor on the back of the robot was positioned against the window, measuring outdoor illuminance from sunlight. Eq. 1 presents the proportional controller used on each SGbot.

$$u_i(t) = k_p e(t) + b \tag{1}$$

$$\text{where} \quad e(t) = x_i(t) - x_i(t-1) \text{ and } \quad x_i(t) = s_i(t)$$

$u_i(t)$ was the control output in the form of Pulse Width Modulation (PWM) signal value given to the motor of SGbot i to buckle or flatten the plate. k_p was a proportional gain. b was a control bias term. $e(t)$ was the error between the normalized desired value of the encoding sensor $x_i(t)$ and the normalized current sensor reading $x_i(t-1)$. To achieve proportionality, the desired value

of the encoding sensor, which determined how far the plate was extended - how bloomed the robot was - was set to equal the normalized value of the ambient light sensor $s_i(t)$. More information on illuminance thresholds and the normalization process can be found in the Supplementary Methods. In this configuration, each SGbot operated independently, relying solely on its own sensor to react based on sunlight exposure.

As can be seen from Fig. 3A, the robots responded well to sunlight, indeed buckling and flattening proportionally to the amount of ambient light sensed. Movie S1 shows a sped up video of the experiment. The Pearson correlation coefficient between the timeseries data throughout the experiment of the encoding sensor values and the ambient light sensor values on that day was $r = 0.98$. Fig. 3B shows the results from a cloudy day (movie S2 shows a sped up video of the experiment), with the Pearson correlation coefficient of $r = 0.95$. The mean, standard deviation and p-value of the correlation coefficients for all three days (of which two were sunny and one was cloudy, see third day results in fig. S1) were: $\mu = 0.963$, $\sigma = 0.015$, $p < 0.0001$.

The high correlation in the results confirms that the robot's response to illuminance was proportional, even under varying weather conditions. However, in this individualistic approach, robots did not take into account any user preferences for desired indoor illuminance, nor did they cooperate with their neighbors to take advantage of redundancy in case of failures. The next study addressed these issues using swarm intelligence.

Adaptive shading using collective opinion dynamics

We aimed to test whether we can enhance the system's robustness and incorporate a broader range of influencing factors. At the same time, we wanted to maintain the desirable proportional response to sunlight obtained from the previous experiments. Therefore, we developed a mathematical model based on opinion dynamics (*51*), called the Swarm Garden opinion dynamics model (SG_od). Opinion dynamics is a well-known swarm intelligence framework, often used to model collective decision-making in natural swarms such as honey bees (*52*). It can be used to enable agents in a swarm to integrate various factors to make a decision, including input from neighboring agents and environmental cues. Similarly, the SG_od model (Eq. 2) allowed each robot to make decisions based on three terms: its own ambient light sensor readings, those of its neighbors and the preferred ambient light level in the room. In contrast to linear consensus models like the DeGroot model (*53*),

the SG_od model introduced non-linearity and adaptive weights. The non-linearity enabled stability by deactivating the model when environmental conditions align with user preferences and activating only when adaptation is needed. Adaptive weights supported sensor fusion by adjusting the influence of each term in response to failures and lighting conditions.

We installed an ambient light sensor in the room to measure its illuminance levels (that is, the indoor illuminance). The sensor's data was then incorporated into the SG_od model as an influencing factor to the overall buckling of SGbots to help regulate the room's lighting based on a user's preference. We assumed the user preference to be at 750 lux - a value that could increase or decrease based on different users. In addition to user preference, we integrated each SGbot's neighbors' opinions into the model as another influencing factor. By doing so, we aimed to incorporate redundancy to mitigate failures or noise arising from individual sensors. We tested the SG_od model for five days in a similar setup to the individualistic approach experiments; 16 SGbots, seven hours each day, from 7:00 to 14:00.

$$x_i(t) = \triangle \left(\underbrace{\alpha s_i(t)}_{1^{\text{st}} \text{ term (own sensor)}} + \underbrace{\beta \sum_{\substack{j=0 \\ j \neq i}}^{n} \frac{s_j(t)}{n}}_{2^{\text{nd}} \text{ term (neighbors' sensors)}} + \underbrace{\gamma k_r}_{3^{\text{rd}} \text{ term (room's sensor)}} \right) - ((\triangle - 1) \cdot (x_i(t-1))) \tag{2}$$

$$\text{and } \triangle = \begin{cases} 0 & \text{if } -0.2 < \tanh(s_r(t)) < 0.2 \text{ (equivalent to 650 lux} < s_r(t) < 850 \text{ lux)} \\ 1 & \text{otherwise} \end{cases}$$

$$\gamma = 0.5 tanh(s_r(t)), \alpha = 0.5(1-\gamma), \beta = 1 - \alpha - \gamma \tag{3}$$

In the SG_od model (Eq. 2), the SGbot i's $x_i(t)$ was influenced by its own sensor, its neighbors' sensors, and the room sensor; the coefficients α, β and γ were weights to each of those terms (Eq. 3). $s_i(t)$ was the illuminance value of SGbot i's own ambient light sensor, $s_j(t)$ was the illuminance value of a neighboring SGbot j's ambient light sensor and n was the total number of SGbots in the

neighborhood of SGbot i. k_r was a constant scaling factor. $s_r(t)$ was the illuminance value of the room's ambient light sensor. All sensor values were normalized.

The $\triangle$ operator was a non-linear activation function that determined whether the model activated or not (Eq. 2). If the illuminance in the room was within a preferred user range, in this case chosen to be in a close range to 750 lux, then the model was deactivated, $\triangle = 0$, and the robot maintained its status quo (denoted by $x_i(t-1)$). If the room's illuminance started deviating from 750 lux, then the model was activated, $\triangle = 1$, to adjust plate displacement; see fig. S2 for a graph of the activation function. The preferred illuminance, 750 lux, could be adjusted by normalizing the input to the hyperbolic tangent function depending on the user's preference. The hyperbolic tangent function was further used with the third term's coefficient, γ, where both were directly proportional (see third term in Eq. 2 and Eq. 3). Consequently, the room's illuminance had a weight that affected the behavior of the system, especially when the room was too dark or too bright. This was because the more the room's illuminance deviated from 750 lux in either direction, the greater the contribution of the third term. However, if the room was within the preferred range, the model deactivated altogether to maintain that ideal status.

Fig. 4 shows the results from a sunny day (Fig. 4A) and a cloudy day (Fig. 4B), see movies S3 and S4 for a sped up video of each of the experiments, respectively. Figs. 4Ai, Bi show the contribution of each term in Eq. 2 to the final output $x_i(t)$ as the day passed. When the room was too dark, the output of the hyperbolic tangent function (in the third term) was negative, influencing the model and bringing the final value lower. This can be seen in Figs. 4Ai, Bi at the beginning and the end of the experiment, when the sun was not shining directly into the room, rendering the room too dark. This resulted in lowering the final value, $x_i(t)$, enabling full - or close to full - buckling to allow the maximum amount of light in. Otherwise, when the room was too bright, the third term was positive, adding to the terms (as observed in the region with stacked terms) to flatten the plates and block light. The model, therefore, allowed us to incorporate user preference without affecting the desirable proportional behavior observed from the individualistic approach.

The distributed nature of swarm systems gives them the potential to be resilient to individual failures - common in large-scale systems - provided the right coordination algorithms are in place. To test if our system was indeed robust to failures, we introduced deliberate failures in simulation. This was done to ensure the model could correctly readjust its coefficients and recover in case of

failure. The simulations were based on real data collected from the days we ran the SG_od model. Specifically, we tested three scenarios: Communication failure where the robot would only rely on its ambient light sensor (first term), own sensor failure where the robot would rely on its neighbors' ambient light sensors (second term) and the room's ambient light sensor (third term) and finally, room sensor failure where the robot would rely on its own ambient light sensor (first term) and its neighbors' ambient light sensors (second term).

In simulation, we observed the change of behavior of one of the robots, robot 5, in these three scenarios and compared them to the normal scenario (SG_od model running with no failures) obtained from the real-life experiment. Since we could not predict how one robot failure would affect the room's ambient light in the second scenario, we assumed it stayed the same, based on real experiment data.

Figs. 4Aii, Bii show the final value, $x_i(t)$, from robot 5 after running the experiment for the three scenarios and against the scenario with no failures. Fig. 4Aii also shows the Pearson correlation coefficient between all of the scenarios for the sunny day (between $r = 0.95$ and $r = 1.00$) and Fig. Bii shows the same for the cloudy day (between $r = 0.83$ and $r = 1.00$). The rest of the three days - see figs. S3, S4, S5 - showed similar correlations and coefficient distributions, with two sunny days and one rainy day. The high similarity between all the scenarios, even in rainy and cloudy days, showed that the system did indeed recover from multiple failure scenarios by discarding the failing terms and adjusting the model to adapt successfully (as can be seen in Figs. 4Aiii, Biii which show the average of the coefficients in each scenario).

Atrium shading in architectural simulation

To demonstrate the capability of the system to scale to large numbers and meet different objectives, we tested a configuration not accessible in our experiments - an atrium with units arranged on a large skylight (see conceptual render in Fig. 5A). In this scenario, we considered tunability in this scaled-up setting, where the objective was to provide selective illuminance - lighting assigned to different sections of a space according to user preference. To investigate this, we conducted a study numerically using Rhino3D/Grasshopper (McNeel) with the plugins Ladybug Tools and Honeybee. With this numerical study, we simulated a south-facing building with dimensions of 20 m depth, 20 m width, and 15 m height, and covered the atrium window in a 19x19 grid, with

each SGbot scaled to approximately 1m (see Fig. 5B.) The weather data used in simulations were obtained from the Typical Meteorological Year 3 (TMY3) weather file for Trenton Mercer County, New Jersey, USA (WMO 724095), constructed by the National Renewable Energy Laboratory from historical observations for 21 June (*60*). More information on the experiment setup can be found in the Supplementary Methods.

In Fig. 5C, we compared the illuminance at 13:00 on June 21 for six shading configurations: open window (no modules), all modules closed, all modules open halfway, all modules maximally open (to the range set in experiments), a horizontal gradient, and a vertical gradient. The results showed that the wide range of continuous deformation of the modules translated to a wide range of illuminance (see illuminance ranges and contrast ratios in Supplementary Methods). As seen by the shifting illuminance gradients in Fig. 5Cv,vi., additional spatial tunability was indeed accessible in a larger system, meaning that the system could accommodate users with different lighting needs in the same space.

In Fig. 5D, we showed an example of targeted lighting for different illuminances for different tasks within the same space at 9:00, 13:00, and 17:00 on June 21. Changes in module configuration reflected changes in sun position, to ultimately preserve the brighter region centered in the north (lower) portion of the work plane, with gradual dimming elsewhere. The maximum illuminance was kept to approximately 3000 lux, which was within the typical bounds of the Useful Daylight Index (*54*). The difference in the illuminance plots between 5Di (9:00) and 5Dii (13:00) was 8.5%, and that between 5Diii (17:00) and 5Dii (13:00) was 10.2%. Although module configurations were adjusted manually in this demonstration, as could be done via gestures according to changing preference, algorithmic adjustment could likely produce more constant illuminance over time. In comparison to traditional shading technologies, from louvered blinds to actuated shading roofs, the Swarm Garden offered independent light targeting in multiple directions, with additional degrees of freedom. This multidirectional spatial tunability could be accessed in vertical windows, horizontal atrium ceilings, or even more complex 3D structures such as pavilions (see Fig. 1H).

Case study 2: Swarm robots for creative expression in interior design

The Swarm Garden was designed to be deployed in human spaces, where its effectiveness depended not only on its technical functionality, but also on how people perceived and interacted with it to enhance their interior space. To evaluate the user interaction aspect of our system, we needed to test whether the system could demonstrate its ability and durability to function in human environments, withstand the challenges posed by crowds where multiple people interacted with it, engage people positively and effectively, including non-experts, and allow people to express themselves and shape their interior spaces through the Swarm Garden.

To answer these questions, we hosted a public exhibition showcasing 36 SGbots of the Swarm Garden that was held on April 9th, 2024 at Princeton University's Lewis Center for the Arts (see Fig. 6A-F). The public exhibition was held for three hours with mixed media (posters, display cases, a television with video) displayed throughout the room to provide more background, technical information, motivation, and context for the project. Co-author Jones performed a live improvisational dance, equipped with a wearable device (see Materials and Methods) we designed midway through the exhibition, allowing visitors to observe the emergent behavior of the Swarm Garden. Over 100 attendees visited and interacted with the Swarm Garden. We collected attendee feedback, gauging sentiment from the public about their experience with Swarm Garden (approved by Princeton University's Institutional Review Board, IRB# 16722).

We first present the two interaction modes which enabled the crowds to both observe and interact with the swarm, showcasing its self-organizing capabilities and ability to engage with multiple people. We then present the qualitative and quantitative findings from the exhibition, which provided valuable insights for evaluating public sentiments regarding the Swarm Garden. We finally explore the system's ability to be used as means of expression, by presenting results from an interview conducted with Jones on her experience dancing with the system using the wearable.

Passive observation of self-organization

In the absence of human interaction, we programmed the Swarm Garden to show patterns of behaviors to create more salient and larger-scale visuals, capturing attendees' attention and highlighting the full range of the system's aesthetic and self-organizing capabilities. Such dynamic yet passive

form of engagement has been shown previously by designers to create beautiful and lively installations that engage crowds (*55*). This mode we created showed various combinations of ordered and randomized buckling and flattening, as well as light color changes (see Fig. 6C and movie S5). Using the swarm concept of morphogen gradients, the changes appeared to flow and dissipate through the entire swarm from one direction to the other. To enable this mode, we developed a localization method - details can be found in Materials and Methods - that allowed SGbots to identify their immediate neighbors and propagate light/blooming accordingly.

Crowd interaction with single SGbots

In order to directly interact with the swarm, as well as support multi-human interactions, we developed a mode that allowed users to selectively flatten or buckle individual SGbots. A user could adjust a single SGbot to create a localized shaded area without affecting the entire swarm. Moreover, this mode facilitated multi-human interaction, as several individuals could simultaneously interact with different SGbots without disrupting each other's experiences. This granular control not only enhanced the system's versatility but also promoted collaborative and personalized interactions with the Swarm Garden. This mode was achieved through the SGbot's interaction sensor, which detected hand gestures to trigger flattening or buckling of the plate (see Fig. 6D-F and movie S6). Details on the implementation of this mode can be found in the Supplementary Methods. Other selective modes we developed, including a transient trail of SGbots' LED following the user as they passed in front of the Swarm Garden, and painting the Swarm Garden by hand picking different LED lights for individual SGbots, are demonstrated in movies S7 and S8, respectively.

Public Exhibition

During the entire three-hour public exhibition, the Swarm Garden operated smoothly without any issues, demonstrating its durability in handling crowds and switching interaction modes. We ran the passive observation mode for the first 30 minutes of the exhibition while attendees were still arriving and familiarizing themselves with the space. The various patterns grabbed attendees' attention who then were eager to start interacting more closely with the swarm. Therefore, we enabled crowd interaction, allowing attendees to manipulate the plate displacement of the SGbots, offering a hands-on understanding of the system's adaptability and responsiveness. This allowed for many

users to interact with the swarm at once. Midway through the exhibition, Jones performed with the swarm. After her performance, we re-enabled crowd interaction so the attendees could continue engaging with the swarm. We then surveyed the attendees, by requesting them to optionally fill out a form asking: "Describe your experience with Swarm Garden in one word". They were also prompted to provide any other feedback on the exhibition in an open-ended text format.

Fifty-seven unique respondents submitted a word or multiple words (up to three) to the word cloud and 21 respondents left additional comments in a general feedback and reflections form. All responses can be found in supplementary data S1 (word cloud) and S2 (feedback). The word cloud, shown in Fig. 6G, immediately showed a general positive sentiment across the largest words, with the most popular words being "cool" (frequency = 9) and "interactive" (frequency = 8), as shown in the distribution in Fig. 6H. To more formally and objectively define the overall audience sentiment, Python Natural Language Toolkit's (NLTK) VADER (Valence Aware Dictionary for Sentiment Reasoning) sentiment analysis tools were utilized to extract sentiment scores based on the words submitted to the word cloud. By manually assigning negative scores to words such as "manmade", "crowded", and "dim" that were otherwise interpreted as neutral by the model, the sentiments could be categorized more accurately. The audience responded with overwhelmingly positive sentiment (95.8% positive sentiments and 4.2% negative sentiments, shown in Fig. 6H). We supported this finding by the strong positive sentiments expressed by the majority of open-ended form responses: 19 of 21 were positive, one was a neutral question, and one was a constructive criticism (see Table 1 for examples).

The Swarm Garden's "floral", "colorful" and "beautiful" appearance left the attendees generally with an "inspiring" feeling that was triggered with their "interactive" experience while "engaging" with the swarm. The Swarm Garden seemed to have succeeded at capturing the attendees' attention, with words frequently appearing such as "captivating" and "mesmerizing", and still leaving them with a "calming" and "fresh" feeling. The highly positive reception and enthusiasm suggested that dynamic and responsive architectural swarm interiors have the potential to contribute to a sense of joy and well-being.

Creative expression through dance

Single interactions allowed individuals direct control over the SGbots in the Swarm Garden. A more holistic interaction, where users influence the swarm's self-organization, introduced a balance between predictability driven by the user control and randomness driven by the swarm's emergent behavior. This dynamic control could start a two-way conversation between user and swarm, creating a unique experience each time, encouraging improvisation, exploration, and creative expression.

To further investigate the Swarm Garden's performative potential, we collaborated with co-author Azariah Jones, a professional dancer with 19 years of experience across various styles, including social dance, ballet, contemporary dance, and jazz funk. Equipped with the wearable, Jones engaged with the Swarm Garden as a dynamic dance partner, allowing it to complement and interact with her movements throughout the performance. Jones engaged with the swarm in improvisational dance twice privately with the rest of the Swarm Garden team (see movie S9 for a performance) and once in front of a public audience at the public exhibition we hosted with the Swarm Garden. We then interviewed Jones on her experience improvising with the Swarm Garden, which allowed us to study and analyze how the interaction influenced her movements, perception of dance partnership, and artistic expression.

We designed a lightweight wrist wearable (see Fig. 7A), motivated by the idea that a dancer could use natural arm movements in the choreography to guide and shape the response of the swarm. The wearable mapped arm movements to different LED color responses, creating a dialogue between dancer and swarm through body movement and light. The direction of arm movement in all three axis was detected through an accelerometer on the Arduino Nicla Vision in the wearable device, and depending on the direction, the LEDs of the SGbots responded differently (see movie S10 for the responses). The y-axis movement was consistent every time, allowing the user to either turn on all the LEDs to green (negative y-axis), or turn off all LEDs (positive y-axis). A movement in the positive x-axis generated a random number between 0 and 35 to pick a random SGbot to propagate the color fuchsia from those SGbots to the whole swarm. A movement in the z-axis picked a random SGbot to propagate the color blue only to the SGbots to the left of the chosen SGbot (negative z-axis) or the color pink only to the SGbots to the right of the chosen SGbot (positive z-axis). In this way, the arm movements corresponded to both predictable and emergent responses from the

swarm, resulting in a different performance effect even if the choreography remains the same. We next discuss results from the interview on how this interaction mode influenced the dancer in an improvisational setting (see Figs. 7B-F for snapshots of the dances), where the dancer could modify the choreography in real-time in response to observations of the swarm.

Embodied interaction. When reflecting on her experience improvising with the Swarm Garden, Jones described their relationship as a "partnership defined by negotiation". Due to the swarm's inherent unpredictability, the roles of a dance leader and follower became ambiguous. Instead, they were both partners influencing one another; instead of exerting direct control, she engaged in a dynamic "push and pull" interaction, adjusting her movements in response to the swarm's reactions and creating movements that would elicit different responses from the swarm. Because sometimes the cues she gave the swarm through her dance movements resulted in an unpredictable behavior, it felt to her as if the swarm was "alive".

Physical aspects. Although the wearable itself was comfortable and did not restrict movement, Jones noted that dancing with the swarm heightened her awareness of her own body. In particular, it brought attention to her arm movements - a part she had often perceived as "passive" in her dancing - which took on a more active role in shaping her dialogue with the swarm. However, physical constraints emerged due to the limitations of the technology. Jones had to remain within specific spatial boundaries to maintain communication between the wearable and the swarm, subtly influencing the way she moved and engaged with the system.

Audience interaction. Jones explained that, during her live performance, although there were moments when she simply went with the flow without focusing on the Swarm Garden's actions, she was also conscious of fully showcasing the system's capabilities - particularly through her arm movements. This awareness inspired her to "search her mind" for movements that could elicit different responses from the swarm. In this way, she not only influenced the swarm's behavior but was also influenced by it, prompting her to explore ways of expressing herself through dance. Through this exploration, Jones developed what she described as a "secret language" with the swarm. Unlike the audience, who remained unaware of which movements triggered specific responses from the Swarm Garden, she understood these subtle cues, deepening her connection to the swarm. This sense of synchronization made her feel "locked in the zone" and as if she and the swarm were "seeing eye to eye." This partnership between dancer and swarm captivated the audience, who later

approached Jones with questions, intrigued by the nature of this partnership.

Artistic expression. Jones described the experience as one that "opened up for me the limits of what expression and interaction can be within dance." She explained that the Swarm Garden expanded her understanding of how dance could exist beyond the traditional stage, allowing technology to foster new forms of partnership and creativity - essentially transforming it into a new kind of "proscenium". Reflecting further, Jones saw the Swarm Garden as an exciting form of performative art; an interactive exhibition that, like an installation in a museum which could be activated by passersbys, could also be activated through the medium of dance.

DISCUSSION

Architectural swarms have the potential to create adaptive, responsive environments that shape indoor spaces and react to outdoor conditions to meet occupants' needs. The Swarm Garden is a proof of concept that leverages swarm robotics and swarm intelligence algorithms to create an architectural façade that is capable of complex self-organized responses to environment and people. The Swarm Garden is envisioned to be deployed in different configurations and diverse spaces, such as vertically on large windows, horizontally in atrium spaces, or in theaters to act as an extra expressive dimension for performative arts.

The Swarm Garden functions as an adaptive shading mechanism, responding proportionally to daylight. The SG_od model enabled self-organization and handled multiple failure modes, such as communication losses and sensor malfunctions. User preference was incorporated by factoring preferred room ambient light into the model. Results showed high correlation between communication loss and room sensor failure scenarios, and less to the normal scenarios, due to the removal of the third term in the model in both scenarios (reflecting removal of user preference). This shows the high effect the user preference can have on the model. Adjusting the coefficient of the third term to higher or lower could allow users to add more or less control on lighting preference in the future. Other factors, such as time of day or temperature, can be added or swapped in the model with adjustable coefficients as well. Another high correlation was between the normal scenario and the robot sensor failure scenario because robots and their neighbors share similar opinions in such a small window. As the swarm expands to hundreds or thousands of robots and is deployed

across bigger windows where sunlight intensity varies, the SG_od model will operate based on the immediate neighborhood of each robot rather than aggregating input from the entire swarm. This localized approach ensures that robots respond to their specific lighting conditions while being robust to failure, as distant neighbors may experience different sunlight intensities due to the sun's movement. However, far neighbors could still play a role in predicting the sun's trajectory, enabling adaptive responses across the entire system.

Crowd interaction with the Swarm Garden provided valuable feedback for refining its design. Attendees particularly appreciated the system's aesthetics and nature-inspired look. Our work spans a spectrum of interaction modalities with varying levels of user agency - passive observation, crowd interaction, and creative expression, with user agency increasing from less to more, respectively. In the future, we aim to enhance user agency through speech- and tablet-based interfaces for greater accessibility and inclusivity (see Supplementary Methods). Learning-based methods may further increase user agency over time. Some feedback from attendees suggested introducing tunable parameters, such as the speed of blooming. Other designs exploring alternative blooming mechanisms could provide enhanced speed and capabilities (see Supplementary Methods). We used LED responses with a fast-paced creative application like dance for its quicker feedback over blooming. However, the contrast between fast-reacting lights and slower blooming could offer another expressive layer worth exploring, such as correlating blooming to the music. Another promising direction is collaborating with choreographers to explore sensor placement on different body parts. As the dancer noted, the wrist was comfortable and expressive for her, but future iterations could explore alternatives such as belts or anklets, that may be better suited to dance forms that involve different body movements. The successful transformation of the Swarm Garden into a creative partner in dance demonstrates the potential for the system to inspire different forms of creative expression. Results from the interview with the dancer reveal how artists can form a "partnership defined by negotiation" with the swarm. This partnership moves beyond simple control to create interactions that are personally meaningful to the artist and captivating to the audience. By reflecting on both audience–robot (through the survey) and artist–robot (through the interview) interactions, our work contributes to, and deepens our understanding of, human–swarm interaction in performative contexts.

As a step toward integrating the Swarm Garden into human spaces, we plan to collaborate with

architects to evaluate the feasibility of long-term deployments in various configurations and places. Although plastic deformation resulting from stress focusing (*56*) is essential to the blooming-via-buckling process, future work could explore the use of more sustainable, resilient materials than those used in our demonstrations. A promising future direction is to modify the sheet's structure using kirigami-inspired cuts (*57*) to lower actuation power. Long-term user studies will also be conducted to assess user adoption over time and identify opportunities to further customize the system to meet user needs and preferences.

We envision a future where the built environment is increasingly inspired by living architectures, creating façades that constantly adapt to their surroundings and occupants. The Swarm Garden offers a glimpse into that future; an architectural swarm that collectively responds to sunlight and human interaction. It further shapes occupants' spaces by being animated with movement, vibrant with colors, and beautiful in appearance, inspiring creativity and expression.

MATERIALS AND METHODS

The SGbot design

Each SGbot's dimensions are 32.7cm × 32.7cm × 18.1cm. The SGbot's plate is made of a 0.127mm thick plastic shimstock sheet and is secured with a top and bottom clamp. A base Printed Circuit Board (PCB) on the back distributes power and connects the sensors, LED strip, and motor controller to an adaptor PCB that houses the Nicla. Further assembly and electronics details can be found in the Supplementary Methods.

The wearable device design

The wearable device, shown in Fig. 7A, is designed to act as a bracelet to be worn around the wrist. It can be fastened with a Velcro strap that can be adjusted depending on different wrist sizes to accommodate diverse users. The wearable consists of a switch to easily turn it on and off, a 3.3V battery, and an Arduino Nicla, all connected through a stripboard. The components are all neatly hosted in a 3D-printed case that the Velcro runs through. The Nicla has an Inertial Measurement Unit (IMU) onboard, which is used to detect the direction of wrist movement.

Localization and neighbor identification

For passive observation and wearable interactions, SGbots's immediate neighbors are tracked using a localization algorithm. Each SGbot has a unique AprilTag on its back, and a wide-angle camera records the 2D positions of all tags. The six nearest neighbors of each SGbot are identified and assigned fixed relative positions for direction-specific messaging. The camera continuously updates each SGbot's neighbor list to reflect changes in swarm configuration or maintain connectivity even if individual SGbots fail. Additional details on the localization process are provided in the Supplementary Methods.

Statistical analysis

We used the Pearson correlation coefficient to compare the ambient-light and encoding sensor timeseries data for both the individualistic and collective adaptive shading experiments, as well as to assess similarity under different failure modes. For all experiments, we report the p-value, mean, and standard deviation across n trials.

Supplementary materials

Supplementary Methods
Figs. S1 to S11
Tables S1 to S6
Movies S1 to S10
Data S1 to S2

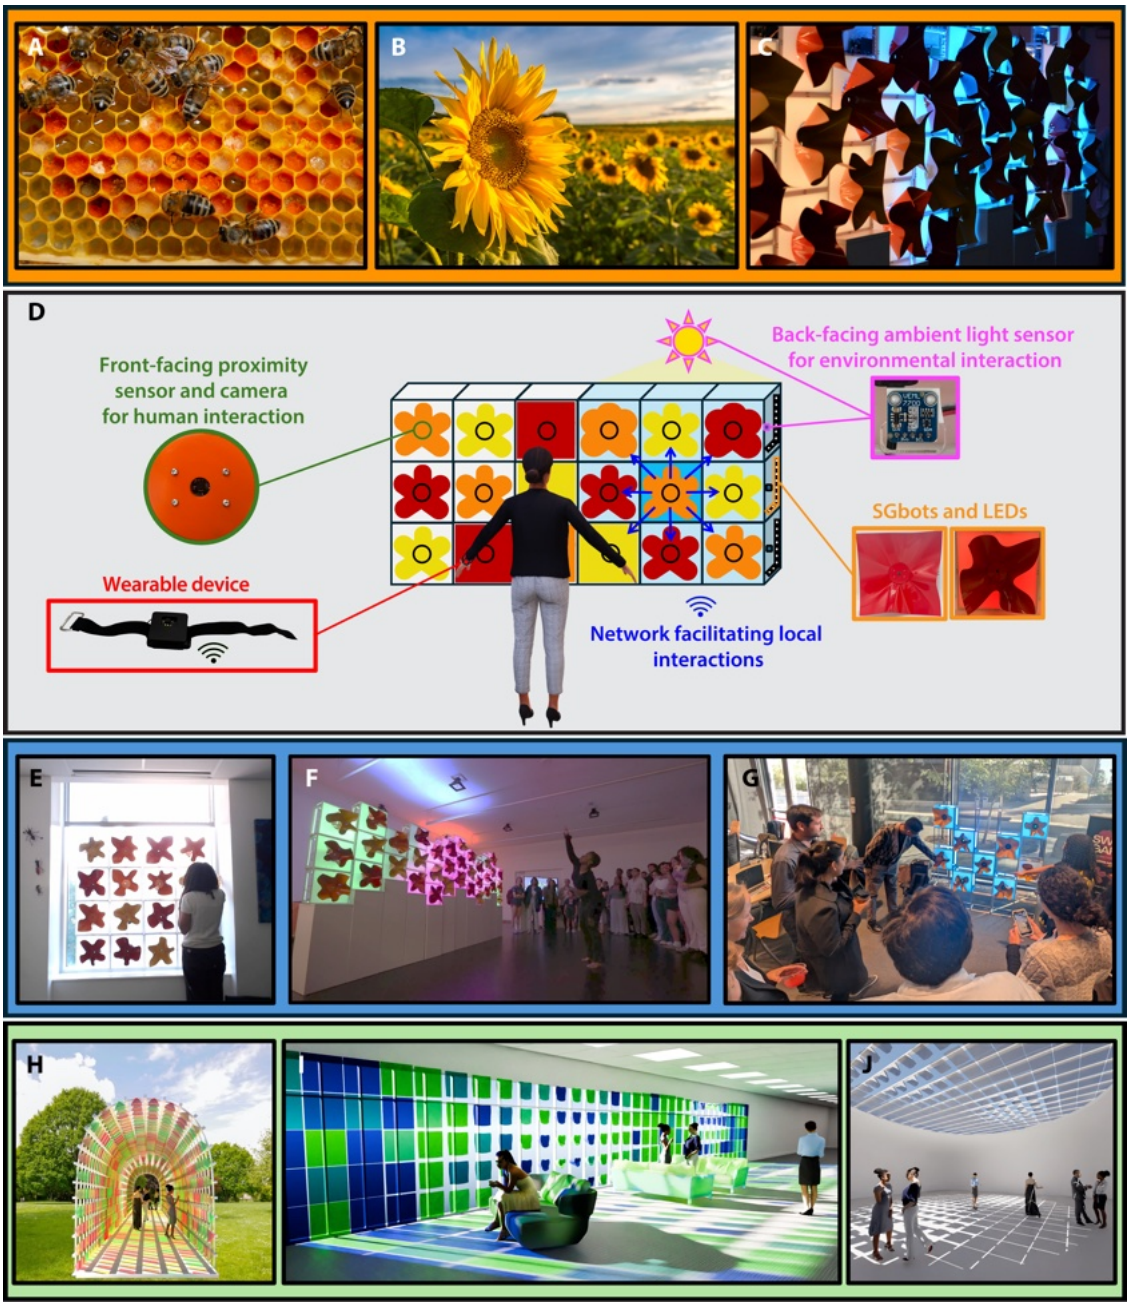

Fig. 1. **The Swarm Garden.** (**A, B**) Bio-inspiration from natural living architectures, like bee hives and fields of flowers, that constantly respond and adapt to their environments. (**C**) SGbots from the Swarm Garden. (**D**) The various components of the Swarm Garden. (**E**) 16 SGbots deployed on an office window for shading. (**F**) A live dance performance with the Swarm Garden during an exhibition. (**G**) The Swarm Garden exhibited as a swarm of sunflowers. (**H-J**) The Swarm Garden envisioned deployed on varied structures. Credits: (A) Shutterstock; (B) Adobe Stock; (D) "Carla Rigged 001 – Rigged 3D Business Women" by Renderpeople (Sketchfab, CC BY 4.0).

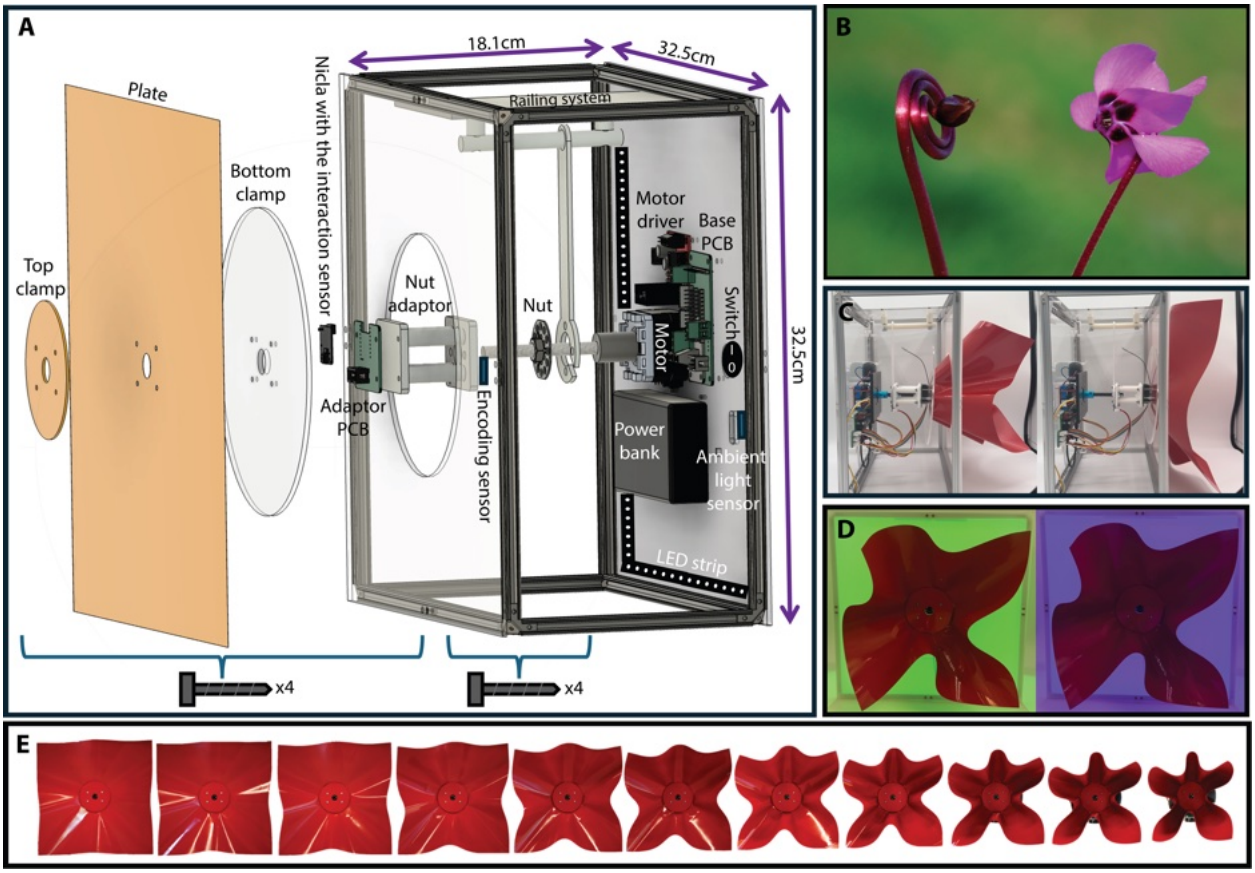

Fig. 2. **The SGbot Design** (**A**) The components and dimensions of the SGbot robot. (**B**) A natural flower blooming (credits: iStock). (**C**) An SGbot blooming. (**D**) Examples showing how frosted acrylic enables light to be diffused across the surface of the acrylic sheet, providing color customization of SGbots. (**E**) The process of blooming; as the plate is pulled through the open ring, the SGbot buckles and contracts.

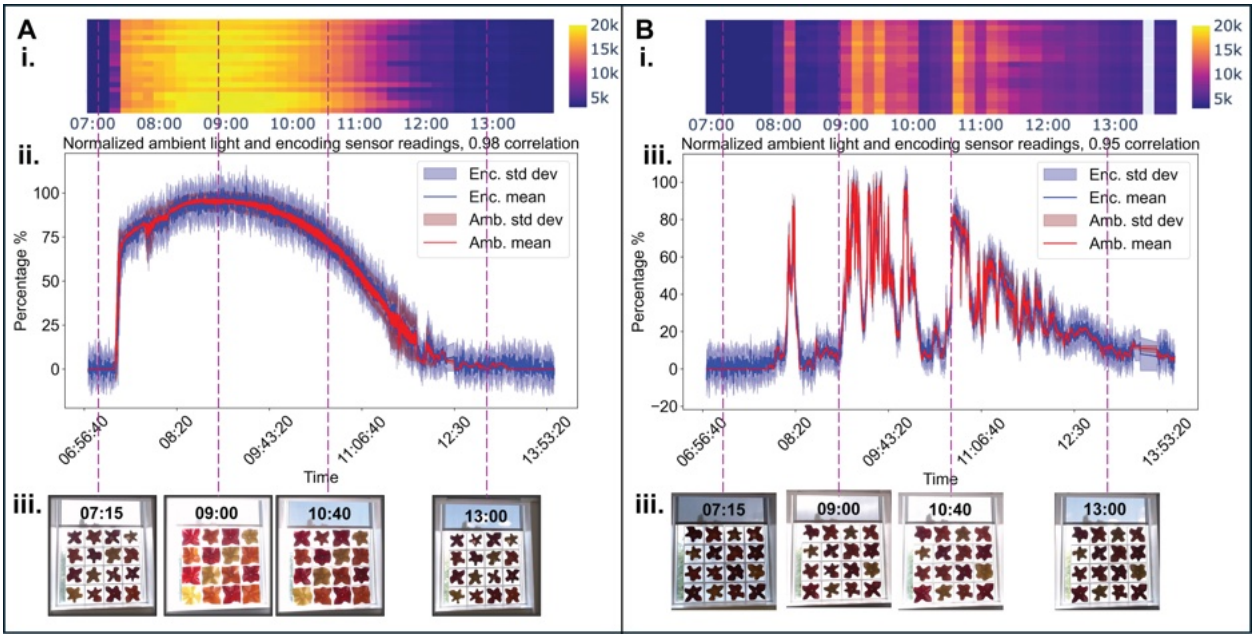

Fig. 3. **Adaptive shading experiment results using an individualistic approach run from 7:00 to 14:00 in a real office.** (**A**) Results on a sunny day (i.) Heatmap showing the raw illuminance values measured by each SGbot during the experiment. Each row in the heat map corresponds to one of 16 SGbots. (ii.) The 10-minute running mean of the normalized encoding sensor readings, corresponding to plate displacement, are shown as a solid blue line averaged across all 16 robots, with the shaded blue region indicating the standard deviation from the mean across all robots (n = 16). The 10-minute running mean of the normalized ambient light sensor readings are shown as a solid red line averaged across all 16 robots, with the shaded red region indicating the standard deviation from the mean across all robots (n = 16).(iii.) Shots of the window shown in different timestamps. (**B**) (i-iii.) Results on a cloudy day running the same experiment.

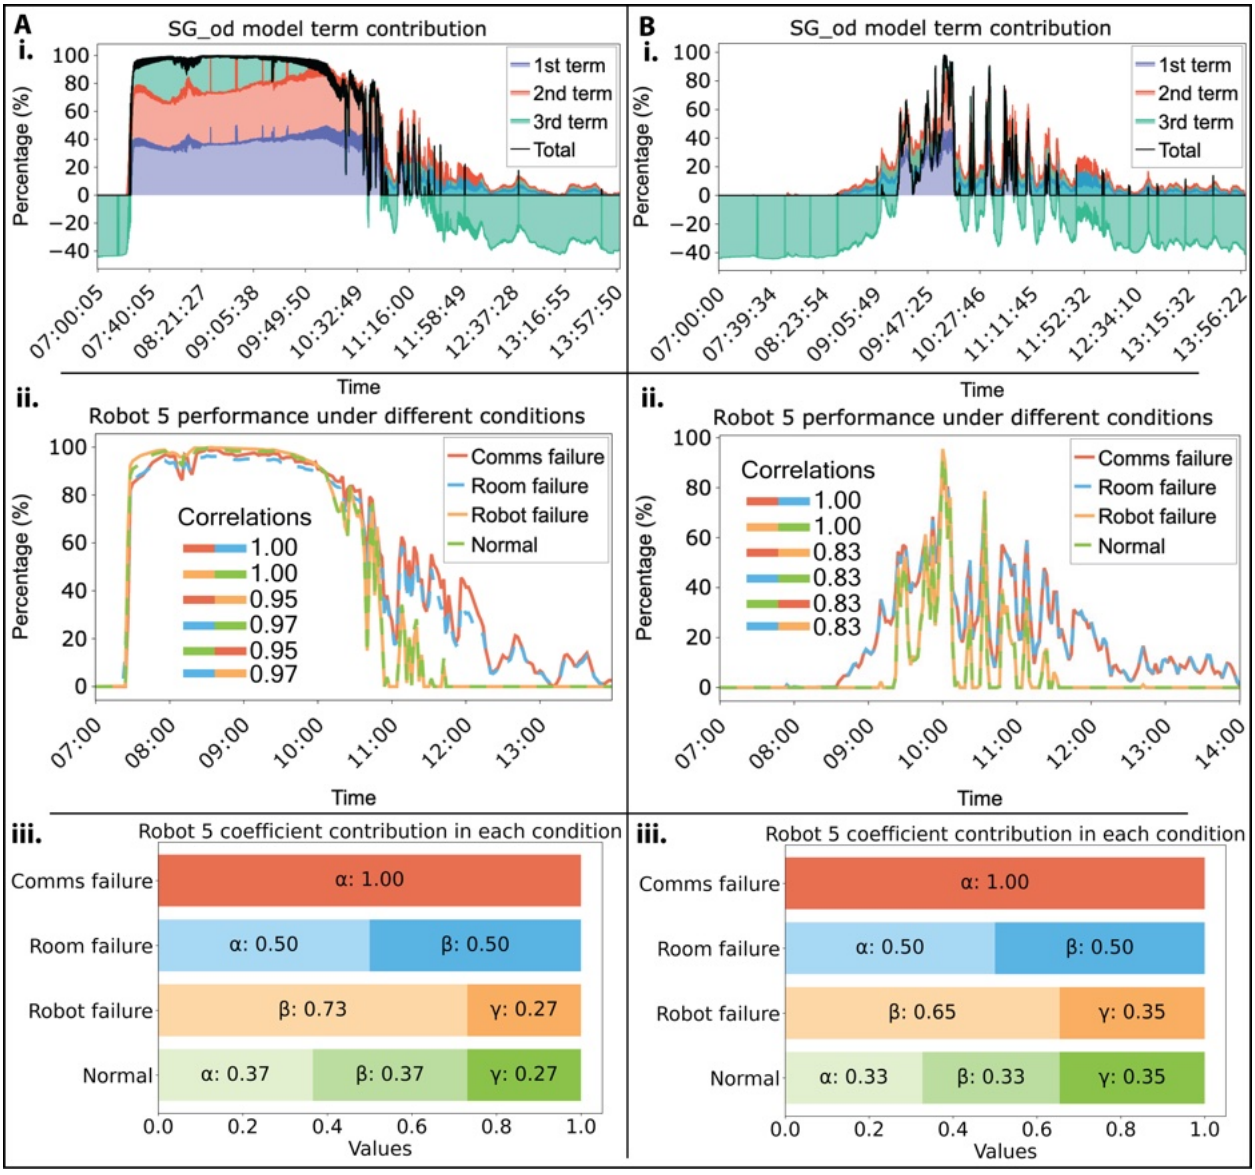

Fig. 4. **Adaptive shading experiment results using collective opinion dynamics run from 7:00 to 14:00 in a real office.** (**A**) Results from a sunny day. (i.) Each term's contribution from the SG_od model (see Eq. 2) represented using a different color, with the final value (total) $x_i(t)$ represented in black. (ii.) Final value, $x_i(t)$, when introducing four failure scenarios in simulation to SGbot 5 for the same day: red is for communication failure, blue is for the room's ambient light sensor failure, orange is for SGbot's ambient light sensor failure and green is for no failures (normal scenario). The Pearson correlation coefficients between each of the 4 scenarios are shown on the figure. (iii.) Average value of each coefficient contribution (used in Eq. 3) in each failure scenario. (**B**) (i-iii.) Results from a cloudy day running the same experiments.

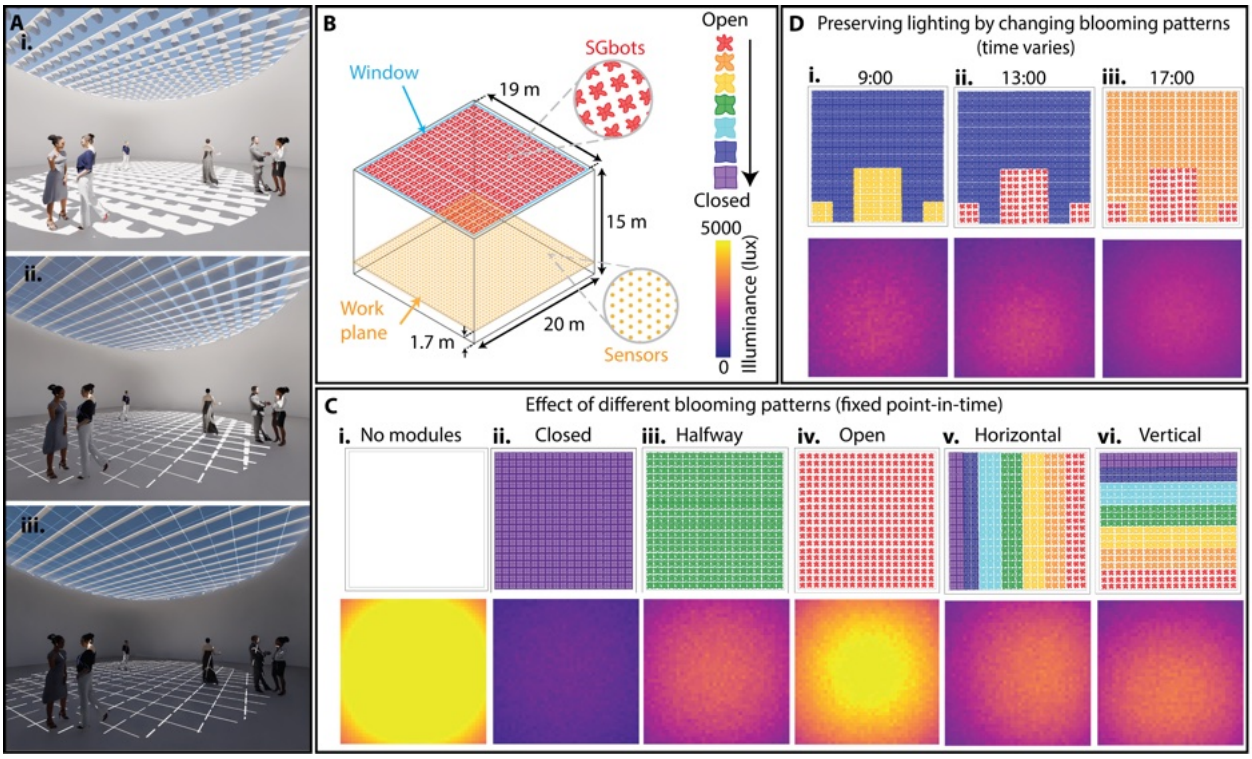

Fig. 5. **Simulated shading configurations in an atrium space.** (**A**) Conceptual render of modules (created using Blender) as a shading system on an atrium skylight window. Modules open in (i.) to allow more light, partially close in (ii.) and close entirely for a dim environment in (iii.) (**B**) Geometry of the simulated building. Dots on the work plane indicate sensor points. Legends on the right corresponds to C and D. (**C**) Skylight module configurations colored by blooming amount as indicated in (B) and resulting illuminance on the work plane colored by illuminance (lux, W/m^2) at 13:00 on June 21 (TMY3) in Mercer County, New Jersey, USA. (**D**) Changing module configurations to preserve similar illuminance conditions throughout the same day, with timestamps indicated above subfigures.

Fig. 6. **The Swarm Garden public exhibition and audience sentiment results.** (**A-B**) The exhibition space. (**C**) Snapshot from the passive observation interaction mode. (**D-F**) Attendees interacting with the Swarm Garden. (**G**) Word cloud of one-word responses from 57 respondents during The Swarm Garden exhibition. Each participant was able to submit up to three unique words to the word cloud. (**H**) Distribution of words appearing more than once in the responses and pie chart showing positive/negative word sentiment.

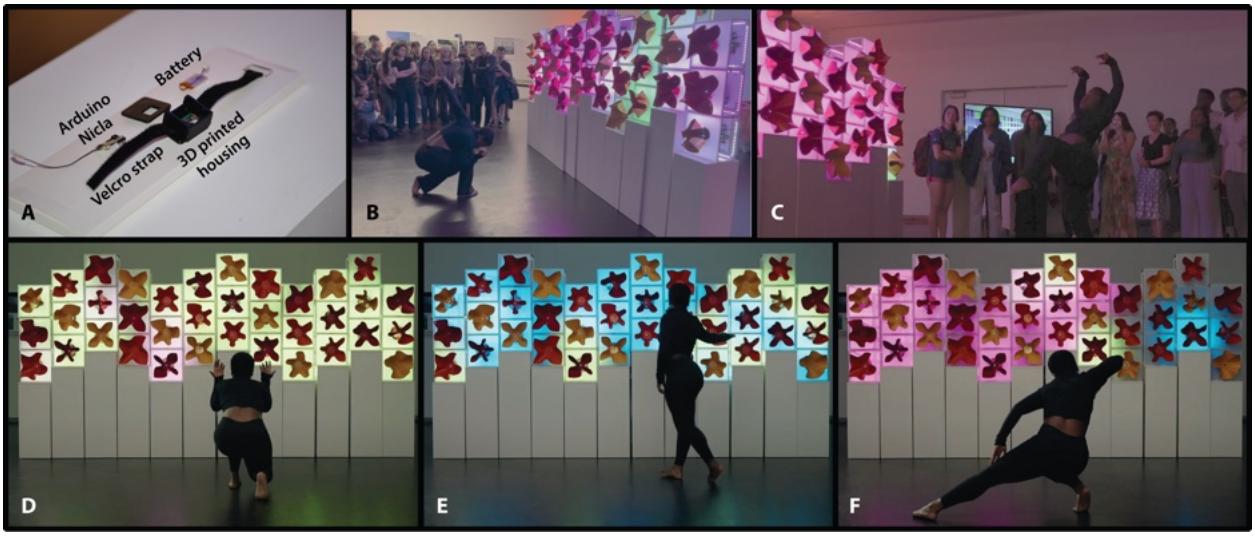

Fig. 7. **Improvisational dance performance with the wearable device.** (**A**) Wearable device components. (**B-C**) Snapshots from the public exhibition improvisational dance performance. (**D-F**) Different movements causing different responses from the Swarm Garden.

Movie 1: **Architectural swarms overview video.** Introduction to architectural swarms and demonstrations of the Swarm Garden's various applications and deployments.

Response type	Quote
Positive	"Such a beautiful and vibrant experience!" "Such a unique display of robotics! The dance interaction is amazing!" "This is sooo cool, it was so amazing to watch and the technology is even crazier!" "I loved how interactive the project is! This is frankly amazing!" "Lovely visuals and great interactive experience"
Neutral	"Can they be programmed to create certain swarm patterns?"
Critiques	"Increase the harmony of the swarm. Make it move more quickly."

Table 1: **Example quotes from the feedback form available to attendees at the Swarm Garden public exhibition.**

References and Notes

1. S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau, *Self-organization in Biological Systems* (Princeton University Press) (2003).

2. V. Charpentier, P. Hannequart, S. Adriaenssens, O. Baverel, E. Viglino, S. Eisenman, Kinematic amplification strategies in plants and engineering. *Smart Materials and Structures* **26** (6), 063002 (2017).

3. H. F. McCreery, Z. A. Dix, M. D. Breed, R. Nagpal, Collective strategy for obstacle navigation during cooperative transport by ants. *Journal of Experimental Biology* **219** (21), 3366–3375 (2016).

4. M. Brambilla, E. Ferrante, M. Birattari, M. Dorigo, Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence* **7** (1), 1–41 (2013).

5. S. Adriaenssens, L. Rhode-Barbarigos, A. Kilian, O. Baverel, V. Charpentier, M. Horner, D. Buzatu, Dialectic Form Finding of Passive and Adaptive Shading Enclosures. *Energies* **7** (8), 5201–5220 (2014).

6. M. Hemberg, U.-M. O'Reilly, A. Menges, K. Jonas, M. d. C. Gonçalves, S. R. Fuchs, *Genr8: Architects' Experience with an Emergent Design Tool* (Springer Berlin Heidelberg, Berlin, Heidelberg), pp. 167–188 (2008).

7. U.-M. O'Reilly, M. Hemberg, Integrating generative growth and evolutionary computation for form exploration. *Genetic Programming and Evolvable Machines* **8** (2), 163–186 (2007).

8. T. Hai, A. El-Shafay, A. Alizadeh, K. Kulshreshtha, S. F. Almojil, A. I. Almohana, A. Alali, Improved locust swarm optimization algorithm applied for building retrofitting based on the green policy of buildings. *Journal of building engineering* **70**, 106274 (2023).

9. N. Delgarm, B. Sajadi, F. Kowsary, S. Delgarm, Multi-objective optimization of the building energy performance: A simulation-based approach by means of particle swarm optimization (PSO). *Applied Energy* **170** (C), 293–303 (2016).

10. M. K. Heinrich, S. Von Mammen, D. N. Hofstadler, M. Wahby, P. Zahadat, T. Skrzypczak, M. D. Soorati, R. Krela, W. Kwiatkowski, T. Schmickl, *et al.*, Constructing living buildings: a review of relevant technologies for a novel application of biohybrid robotics. *Journal of the Royal Society Interface* **16** (156), 20190238 (2019).

11. A. Çelekli, İ. Yeşildağ, Ö. E. Zariç, Green building future: algal application technology. *Journal of Sustainable Construction Materials and Technologies* **9** (2), 1 (2024).

12. M. Wahby, M. K. Heinrich, D. N. Hofstadler, E. Neufeld, I. Kuksin, P. Zahadat, T. Schmickl, P. Ayres, H. Hamann, Autonomously shaping natural climbing plants: a bio-hybrid approach. *R. Soc. Open Sci.* **5** (10), 180296 (2018).

13. M. Brzezicki, A Systematic Review of the Most Recent Concepts in Kinetic Shading Systems with a Focus on Biomimetics: A Motion/Deformation Analysis. *Sustainability* **16** (13), 5697 (2024).

14. N. Heidari Matin, A. Eydgahi, Technologies used in responsive facade systems: a comparative study. *Intelligent Buildings International* **14** (1), 54–73 (2019).

15. F. Sommese, L. Badarnah, G. Ausiello, A critical review of biomimetic building envelopes: towards a bio-adaptive model from nature to architecture. *Renewable and Sustainable Energy Reviews* **169**, 112850 (2022).

16. A. M. A. Faragalla, S. Asadi, Biomimetic Design for Adaptive Building Façades: A Paradigm Shift towards Environmentally Conscious Architecture. *Energies* **15** (15), 5390 (2022).

17. E. Vazquez, D. Correa, S. Poppinga, A review of and taxonomy for elastic kinetic building envelopes. *Journal of Building Engineering* **82**, 108227 (2024).

18. M. Gonçalves, A. Figueiredo, R. Almeida, R. Vicente, Dynamic façades in buildings: A systematic review across thermal comfort, energy efficiency and daylight performance. *Renewable and Sustainable Energy Reviews* **199**, 114474 (2024).

19. M. Pesenti, G. Masera, F. Fiorito, Exploration of Adaptive Origami Shading Concepts through Integrated Dynamic Simulations. *Journal of Architectural Engineering* **24** (4) (2018).

20. L. Le-Thanh, T. Le-Duc, H. Ngo-Minh, Q.-H. Nguyen, H. Nguyen-Xuan, Optimal design of an Origami-inspired kinetic façade by balancing composite motion optimization for improving daylight performance and energy efficiency. *Energy* **219**, 119557 (2021).

21. M. Meloni, J. Cai, Q. Zhang, D. Sang-Hoon Lee, M. Li, R. Ma, T. E. Parashkevov, J. Feng, Engineering Origami: A Comprehensive Review of Recent Applications, Design Methods, and Tools. *Advanced Science* **8** (13) (2021).

22. D. Misseroni, P. P. Pratapa, K. Liu, B. Kresling, Y. Chen, C. Daraio, G. H. Paulino, Origami engineering. *Nature Reviews Methods Primers* **4** (1) (2024).

23. Y. Li, Y. Zhao, Y. Chi, Y. Hong, J. Yin, Shape-morphing materials and structures for energy-efficient building envelopes. *Materials Today Energy* **22**, 100874 (2021).

24. T. Cheng, Y. Tahouni, E. S. Sahin, K. Ulrich, S. Lajewski, C. Bonten, D. Wood, J. Rühe, T. Speck, A. Menges, Weather-responsive adaptive shading through biobased and bioinspired hygromorphic 4D-printing. *Nature Communications* **15** (1) (2024).

25. D. Sung, Smart Geometries for Smart Materials: Taming Thermobimetals to Behave. *Journal of Architectural Education* **70** (1), 96–106 (2016).

26. M. Dorigo, G. Theraulaz, V. Trianni, Swarm Robotics: Past, Present, and Future [Point of View]. *Proceedings of the IEEE* **109** (7), 1152–1165 (2021).

27. K. H. Petersen, N. Napp, R. Stuart-Smith, D. Rus, M. Kovac, A review of collective robotic construction. *Science Robotics* **4** (28), eaau8479 (2019).

28. F. Augugliaro, S. Lupashin, M. Hamer, C. Male, M. Hehn, M. W. Mueller, J. S. Willmann, F. Gramazio, M. Kohler, R. D'Andrea, The Flight Assembled Architecture installation: Cooperative construction with flying machines. *IEEE Control Systems Magazine* **34** (4), 46–64 (2014).

29. S. Parascho, I. X. Han, S. Walker, A. Beghini, E. P. G. Bruun, S. Adriaenssens, Robotic vault: a cooperative robotic assembly method for brick vault construction. *Construction Robotics* **4** (3), 117–126 (2020).

30. J. Werfel, K. Petersen, R. Nagpal, Designing Collective Behavior in a Termite-Inspired Robot Construction Team. *Science* **343** (6172), 754–758 (2014).

31. S. Hauser, M. Mutlu, P.-A. Léziart, H. Khodr, A. Bernardino, A. Ijspeert, Roombots extended: Challenges in the next generation of self-reconfigurable modular robots and their application in adaptive and assistive furniture. *Robotics and Autonomous Systems* **127**, 103467 (2020).

32. L. Tang, D. Rüegg, R. Zhang, A. Bolotnikova, J. Rabaey, A. Ijspeert, Velocity Potential Field Modulation for Dense Coordination of Polytopic Swarms and Its Application to Assistive Robotic Furniture. *IEEE Robotics and Automation Letters* **10** (7), 7452–7459 (2025).

33. F. Mondada, M. Bonani, F. Riedo, M. Briod, L. Pereyre, P. Retornaz, S. Magnenat, Bringing Robotics to Formal Education: The Thymio Open-Source Hardware Robot. *IEEE Robotics & Automation Magazine* **24** (1), 77–85 (2017).

34. A. Ozgur, S. Lemaignan, W. Johal, M. Beltran, M. Briod, L. Pereyre, F. Mondada, P. Dillenbourg, Cellulo: Versatile Handheld Robots for Education, in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '17 (Association for Computing Machinery, New York, NY, USA) (2017), p. 119–127.

35. M. Alhafnawi, S. Hauert, P. O'Dowd, Robotic Canvas: Interactive Painting onto Robot Swarms. *Artificial Life Conference Proceedings* (32), 163–170 (2020).

36. M. Santos, G. Notomista, S. Mayya, M. Egerstedt, Interactive Multi-Robot Painting Through Colored Motion Trails. *Frontiers in Robotics and AI* **7** (2020).

37. N. Correll, N. Farrow, S. Ma, HoneyComb: a platform for computational robotic materials, in *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*, TEI '13 (Association for Computing Machinery, New York, NY, USA) (2013), p. 419–422.

38. M. Alhafnawi, E. R. Hunt, S. Lemaignan, P. O'Dowd, S. Hauert, MOSAIX: a Swarm of Robot Tiles for Social Human-Swarm Interaction, in *2022 International Conference on Robotics and Automation (ICRA)* (2022), pp. 6882–6888.

39. M. Le Goc, L. H. Kim, A. Parsaei, J.-D. Fekete, P. Dragicevic, S. Follmer, Zooids: Building Blocks for Swarm User Interfaces, in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16 (Association for Computing Machinery, New York, NY, USA) (2016), p. 97–109.

40. F. Meggers, D. Aviv, V. Charpentier, E. Teitelbaum, A. Ainslie, S. Adriaenssens, Co-optimization of Solar Tracking for Shading and Photovoltaic Energy Conversion, in *Proceedings of Building Simulation 2017: 15th Conference of IBPSA* (IBPSA, San Francisco, USA), vol. 15 of *Building Simulation* (2017), pp. 2224–2231.

41. V. Charpentier, F. Meggers, S. Adriaenssens, O. Baverel, Occupant-centered optimization framework to evaluate and design new dynamic shading typologies. *PLOS ONE* **15** (4), e0231554 (2020).

42. T. Solano, A. Bernal, D. Mora, M. Chen Austin, How bio-inspired solutions have influenced the built environment design in hot and humid climates. *Frontiers in Built Environment* **9** (2023).

43. Y. Zeng, H. Sun, B. Lin, Q. Zhang, Non-visual effects of office light environment: Field evaluation, model comparison, and spectral analysis. *Building and Environment* **197**, 107859 (2021).

44. P. Zumthor, Atmospheres: Architectural environments. surrounding objects, in *Atmospheres* (Birkhäuser) (2006).

45. R. Scruton, *The aesthetics of architecture* (Princeton University Press) (2021).

46. A. LaViers, Robots and dance: A promising young alchemy. *Annual Review of Control, Robotics, and Autonomous Systems* **8** (1), 323–350 (2025).

47. L. B. Rosselló, M. Alkilabi, E. Tuci, H. Bersini, A. Reina, Emergent Orchestras: A Modular Framework for Musical Robot Swarms, vol. ALIFE 2024: Proceedings of the 2024 Artificial Life Conference of *Artificial Life Conference Proceedings* (2024), pp. 29–37.

48. D. St-Onge, U. Côté-Allard, K. Glette, B. Gosselin, G. Beltrame, Engaging with Robotic Swarms: Commands from Expressive Motion. *J. Hum.-Robot Interact.* **8** (2) (2019).

49. L. Stein-Montalvo, A. Guerra, K. Almeida, O. Kodio, D. P. Holmes, Wrinkling and developable cones in centrally confined sheets. *Physical Review E* **108** (3) (2023).

50. K. A. Seffen, Truncated cones from indenting a clamped disk. *Physical Review E* **110** (3) (2024).

51. N. Leonard, A. Bizyaeva, A. Franci, Fast and Flexible Multiagent Decision-Making. *Annual Review of Control, Robotics, and Autonomous Systems* **7** (2023).

52. R. Gray, A. Franci, V. Srivastava, N. E. Leonard, Multiagent Decision-Making Dynamics Inspired by Honeybees. *IEEE Transactions on Control of Network Systems* **5** (2), 793–806 (2018).

53. A. Bizyaeva, A. Franci, N. E. Leonard, Sustained oscillations in multi-topic belief dynamics over signed networks, in *2023 American Control Conference (ACC)* (2023), pp. 4296–4301.

54. J. Mardaljevic, M. Andersen, N. Roy, J. Christoffersen, Daylighting metrics: is there a relation between useful daylight illuminance and daylight glare probabilty?, in *Proceedings of the building simulation and optimization conference BSO12* (2012), pp. 189–196.

55. M. Salimi, *Swarm Systems in Art and Architecture: State of the Art* (Springer Nature) (2021).

56. T. A. Witten, Stress focusing in elastic sheets. *Rev. Mod. Phys.* **79** (2), 643–675 (2007).

57. L. Stein-Montalvo, L. Ding, M. Hultmark, S. Adriaenssens, E. Bou-Zeid, Kirigami-inspired wind steering for natural ventilation. *Journal of Wind Engineering and Industrial Aerodynamics* **246**, 105667 (2024).

58. C. Hutto, E. Gilbert, VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media* **8** (1), 216–225 (2014).

59. J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in *North American Chapter of the Association for Computational Linguistics* (2019).

60. S. Wilcox, W. Marion, Users manual for tMY3 data sets (revised) (no. nRel/tP-581-43156,national Renewable energy lab, 2008).

Acknowledgments

We thank Dr. Amy LaViers for her insightful comments on our paper. We thank Sneha Ramshanker for proofreading and providing editing suggestions. We thank the SSR Lab members: Dr. Hungtang Ko, Dr. Di Ni, Irene Sha and Sneha Ramshanker for their help in setting up exhibitions. We thank Baffour Osei for providing us with technical expertise. Inspired by a recent initiative to increase awareness and mitigate citation bias (Zurn et al https://doi.org/10.1016/j.tics.2020.06.009), we include a gender citation diversity statement using manually compiled data. Our references contain 15.3% woman (first author)/woman (last author), 8.5% man/woman, 22% woman/man, 50.8% man/man and 3.4% unknown. Due to limitations on prediction of other marginalizations, we do not yet include other statistics but aim to do so in the future. We look forward to future work that could help us better understand how to support equitable practices in science.

Funding: This work was supported by the National Science Foundation grant CMMI-2240407, CreativeX and Amazon Research Awards.

Author contributions: R.N. and S.A. conceptualized and supervised the project. M.A. led the development of the project, designed the hardware of the SGbot, designed the algorithms for and carried out the office shading experiments, and acquired IRB approval. L.S.M carried out the atrium numerical experiments and introduced the sheet confinement concept. J.B., Y.T. and M.A. scaled the robot hardware into a swarm. J.B. and Y.T. developed the algorithms for the human interaction modes, the wearable interactions and the localization algorithms. Y.T. and M.A. designed the wearable device. A.J., Y.T. and J.B. developed the performative piece with the robots. M.A. analyzed the dance interview results and J.B, Y.T and M.A. analyzed the public exhibition results. V.C. developed renderings and exhibition pieces. The entire team visualized and carried out the public exhibition.

Competing interests: There are no competing interests to declare.

Data, Code, Materials Availability: All data needed to evaluate the conclusions of this manuscript are present in the main text, the supplementary materials, or the Dryad repository (https://doi.org/10.5061/dryad.5x69p8djj) which contains shading experiments data and code, as well as the interaction code.

Supplementary materials

Supplementary Methods

Figs. S1 to S11

Tables S1 to S6

Movies S1 to S10

Data S1 to S2

Supplementary Materials for

Architectural swarms for responsive façades and creative expression

Merihan Alhafnawi[1*], Jad Bendarkawi[2†], Yenet Tafesse[3†],Lucia Stein-Montalvo[4], Azariah Jones[5], Vicky Chow[6],Sigrid Adriaenssens[7], Radhika Nagpal[1,3]

[1]Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[2]Department of Electrical and Computer Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[3]Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA.

[4]Department of Civil and Environmental Engineering, Northwestern University, Evanston, Illinois 60208, USA.

[5]Department of African American Studies, Princeton University, Princeton, New Jersey 08544, USA.

[6]School of Architecture, Princeton University, Princeton, New Jersey 08544, USA.

[7]Department of Civil and Environmental Engineering, Princeton University, Princeton, New Jersey 08544, USA.

[*]Corresponding author. Email: m.alhafnawi@princeton.edu

[†]These authors contributed equally to this work.

This PDF file includes:

Other Supplementary Materials for this manuscript:

Supplementary Methods

The Swarm Garden system design: SGbot

In this section, we provide more detailed technical descriptions of the SGBot. Specifically, we add more information on the SGbot hardware design, an SGbot's bill of materials (see table S1), various SGbot control latencies (see table S2), and the SGbot's battery life for different modes (see table S3).

Each SGbot's dimensions are 32.5cm by 32.5cm by 18.1cm. The plate of the SGbot is made of a 0.127mm thick plastic shimstock sheet which is spray painted. The bottom clamp, in Fig. 2A, and all the components behind it are located behind the front-side acrylic sheet (as can be seen in Fig. 2C), whereas the plate and the top clamp are located in front of the front-side acrylic sheet. The bottom clamp has a diameter of 19cm, which is larger than the circular opening (which has a diameter of 15cm) to prevent the mechanism from falling off the rod in case of motor or encoding sensor failures. The top clamp, plate, bottom clamp, adaptor PCB which hosts the Nicla, and the top part of the nut adaptor are all connected with four screws, as indicated by the curly bracket on the bottom left of Fig. 2A. The bottom part of the nut adaptor, the nut, and the wrench-like vertical part of the railing system, are all connected also via four screws, as indicated by the curly bracket on the bottom right of Fig. 2A. The railing system, which consists of a 3D-printed rail, a wrench-like vertical piece that connects the rail to the nut, and an acrylic sheet which the rail is mounted onto, ensures the movement is strictly horizontal (in the z axis), preventing the sheets from rotating off-axis.

A base PCB of our design, mounted on the back-side acrylic sheet of the robot, hosts the circuitry that connects the sensors and motor driver to an Input/Output (I/O) expander. The PCB also connects a power bank to the motor driver, Nicla, the robot's on/off switch, the sensors (ambient light and encoding) and the LED strip to provide power. The LED strip, power supply, and I/O expander are connected to an adaptor PCB housing the Nicla via a ribbon cable.

During initial testing, actuator failures occasionally occurred through missed stepper counts, but adding encoder-based closed-loop control resolved these issues. Networking failures were rarely observed, and the interaction sensor was generally reliable. Minor sensor conflicts between the ambient light and encoder sensors on the shared Inter-Integrated Circuit (I^2C) bus were mitigated with

retry logic, improving robustness. Additionally, decentralized cooperation via opinion dynamics further reduced the effect of ambient light sensor failures.

The SGbot's plate is actuated using a stepper motor which is coupled to a threaded rod. The rod retracts and extends the plate. We also considered alternative actuation mechanisms, particularly with respect to reliability, speed, and precision. We discuss three options: Shape Memory Alloys (SMAs), a pulley-based actuation system and pneumatic actuation.

Shape Memory Alloys (SMAs) contract when heated by electric current and passively return to their original shape upon cooling. However, achieving precise intermediate states (for example, 80% blooming) is difficult due to their nonlinear, temperature-dependent behavior. Additionally, because our robots may be exposed to direct sunlight, external heating could interfere with the SMA response, further reducing precision and reliability. SMAs are also relatively slow in transitioning between states.

A pulley-based actuation system using the same motor as our current threaded rod system was also explored. Although precision and speed would likely remain comparable to the current design, this approach would introduce additional moving parts prone to jamming and alignment issues, ultimately reducing system reliability.

Pneumatic actuation offers fast response times but would require bulky infrastructure–such as compressors, valves, and tubing–which could compromise the modularity and scalability of the swarm. Furthermore, pneumatic systems are often difficult to control precisely, particularly when aiming to hold or achieve partial deformation states. Air leaks are also common, further decreasing system reliability.

An exciting direction for future work involves modifying the structure of the sheet itself–such as through kirigami-inspired cuts (*57*)–to reduce the power required for actuation. Kirigami designs are easy to fabricate using laser-cutting techniques and can introduce flexibility while preserving the simplicity of single-degree-of-freedom actuation that is present with our design. This could substantially lower power consumption while enabling complex shape transformations.

The Swarm Garden system design: Localization and neighbor identification

For the passive observation mode and wearable device interactions we created with the Swarm Garden, SGbots needed to be aware of their immediate neighbors, and so we created a localization

algorithm. Each SGbot is equipped with a unique AprilTag from the 36H11 tag family, which is attached to its back and referred to as the SGbot ID. A wide-angled localization camera is positioned to capture the back of the swarm, recording the 2D positions of these tags and identifying the six nearest neighboring tags by Euclidean distance (see fig. S11 showing the setup). These neighbors are assigned specific relative positions (top, top right, bottom left, and so forth) to enable direction-specific messaging. A command is continuously sent from the localization camera to ensure that each SGbot's list of neighbors is dynamically updated in response to any configuration changes, supporting a flexible and responsive swarm system. This method helps mitigate the risk of single points of failure from individual SGbot outages, as the swarm will continuously update neighbor assignments across all SGbots, maintaining consistent swarm-level connectivity. Note that for the shading experiments, the robots collected opinions from all neighbors directly, as the window size was sufficiently small, eliminating the need for knowing the neighbor's precise location.

Case study 1: Normalization process and illuminance thresholds

For the individualistic approach and collective opinion dynamics shading experiments, the encoding and ambient light sensors' values are both normalized before being used in the robot controller. This is done as both values are compared to one another proportionally since they correspond to one another. 0% represents complete buckling and low outdoor illuminance (3000 lux or less), and 100% represents complete flattening and high outdoor illuminance (20000 lux or above). Since the room where the experiments took place and its window are relatively small, an illuminance value detected by the SGbot's ambient light sensor–which would be facing the window hence measuring the outdoor illuminance since it is located on the back-side of the SGbot–of less than 3000 lux renders the room too dark, and so the plate would fully buckle at illuminance of 3000 lux or lower (corresponding to 0%). Any value at or above 20000 lux renders the room too bright, and would cause the plate to fully flatten (corresponding to 100%).

Case study 1: Swarm Garden's response to sunlight with the SG_od model

The SG_od model enabled self-organization and was robust to multiple failures, including communication losses and sensor malfunctions, and also incorporating user preferences (see table S4 for a comparison with the individualistic approach). We executed a deeper analysis into the performance

of the SGbots' response to sunlight under the SG_od model. The target range for room illuminance was 650 lux to 850 lux. In practice, there were periods when the room was outside this range due to environmental factors beyond the system's control. For example, when the sun had not yet reached the window in the early morning or when it was later blocked by surrounding buildings, the result was insufficient light, rendering the room too dark. In this scenario (which we call the "Low" period), since we cannot generate light, the robots responded by completely buckling to let in as much light as possible, which is the correct behavior (in other words, room ambient light <650 lux and blooming <10%). On the other hand, when the sun was directly angled towards the window, some light still shone through parts uncovered by the robots, rendering the room too bright. In this scenario (which we call the "High" period), the robots responded by completely flattening to block as much light as possible, which is also the correct behavior (in other words, room ambient light >850 lux and blooming >90%). These extreme conditions (High and Low) together constitute the majority of the sunny days: 78.5% on the first day (fig. S6), 73.8% on the fourth day (fig. S7), and 77.4% on the fifth day (fig. S8). During these periods, the robots were consistently fully open or fully closed, which is the correct behavior.

During the remaining times of the day (which we call the "Middle" period), the system adapted continuously to keep the room as close to the desired range as possible. In this condition, we see the robots reacting to bring the room close to that narrow-preferred user range. On sunny days, the room was exactly in range 73.2% of the time (with a deviation from the range by 22.2%) on the first day; 51.6% (deviation: 11.5%) on the fourth day; and 39.6% (deviation: 27.2%) on the fifth day. The sun and clouds change dramatically on a daily basis, and that is why we see different numbers for each day. We further investigated the cases where the room was not in the desired range. We found that these cases were clustered at the "edges", where the sun is just entering or exiting the window region and therefore the level of sunlight is changing rapidly. This suggests that there is some delay as the system adjusts, which may be improved through further tuning or learning of weights. This also varies between days, for example on the fifth day the period of stable sunlight is much shorter than first day. Even so, the deviation from the range always remained only below 30% on sunny days, suggesting we were always close. On rainy and cloudy days, the system primarily remained in the Low period due to the lack of natural light; the room stayed in the Low period 84.8% on the second day (fig. S9), and 84.1% on the third day (fig. S10), where the robots responded correctly by

buckling to let in as much light as possible. For the Middle periods, the room stayed in the preferred range 26.6% with a deviation of 38.5% on the second day, and 49.5% and deviation of 29.7% on the third day.

Future work could explore learning-based methods to adaptively tune the coefficients of each control term over time, helping the system maintain the ideal range throughout the day across different room configurations and deployment locations. Additionally, integrating artificial lighting–such as the robots' own LEDs–could supplement natural light when it is too dark, enabling the system to more consistently achieve the user-preferred illuminance.

Case study 1: Atrium shading experiment setup

The horizontal window used in the simulation is 19m by 19m, centered on the top surface of the building, and a work plane–in other words, where illuminance sensors are located–spans the planar area at an average human height of 1.7m. The illuminance on the work plane is calculated using the simulation software Honeybee Radiance, with 1600 lighting sensors arranged in a grid dividing the plane into squares. The sensor locations are shown as dots in Fig. 5B. To simulate the Swarm Garden building façade system, photographs of a single module through a range of displacements (see Fig. 2E) were used to create 2D shading surface geometries, which were then imported to Rhino3D, scaled to approximately 1m each (note the scale difference compared to experiments), and tiled to cover the atrium window in a 19m by 19m grid. Calculations of illuminance at each sensor point are performed using weather data for Mercer County, New Jersey, USA (latitude 40.3573 degrees N) from the typical meteorological year (TMY3) dataset. Although our simulations were simplified to 2D projected geometries, incorporating reflections from 3D curved surfaces and optimizing material properties could enable more precise light tailoring.

Case study 1: Atrium study additional analysis

We compute the spatial contrast ratio, defined as the ratio between maximum and minimum illuminance on the work plane, for each simulated configuration (see table S5). The modular, adaptive swarm shading system can reduce glare by avoiding abrupt light transitions while still supporting a range of lighting conditions for varied occupant tasks.

We further compare the scenarios presented in Fig. 5D to static shading. Specifically, we

evaluate the configuration from each timestamp against a static baseline in which there are no modules (empty window as in Fig. 5Ci) or static open modules (fully buckled as in Fig. 5Civ). We compute the average percentage difference between the points on the grid of the two scenarios as shown in Eq. S1.

$$\text{Average Percentage Difference} = \frac{1}{1600}\sum_{i=1}^{1600}\frac{\left|\text{adaptive}_i - \text{static}_i\right|}{\frac{\text{adaptive}_i+\text{static}_i}{2}} \times 100 \tag{S1}$$

We see that the average percentage difference is substantial for no modules (121%, 108% and 53% at 9:00, 13:00 and 17:00 respectively) and also for open modules (96%, 81% and 19%), see table S6. The adaptive shading system not only helps stabilize indoor illuminance throughout the day in spite of the large variation in daylight, but also provides substantially more effective shading.

Case study 2: Crowd interaction implementation

On initialization to crowd interaction with single SGbots mode, each module lights up with the LED color corresponding to its plate color which remains constant for the duration of this mode. When a user is close enough in proximity, the module either buckles or flattens. The motor changes direction once a buckling/flattening threshold is reached to allow users to continuously interact with the modules and observe the plate as it transforms between the two states. Users can then create customized patterns or shapes with the entire swarm by selectively adjusting different SGbots. Each SGbot requires approximately 10 seconds to complete the flattening process and an additional 10 seconds for buckling, resulting in a total blooming time of 20 seconds. The speed of the SGbot is approximately 5 mm/s.

To analyze attendee responses from the exhibition, we utilized VADER sentiment analysis due to its suitability for analyzing short, informal texts, such as the word cloud responses. Developed specifically for social media and conversational language, VADER excels at handling informal expressions and slang (*58*), which were more prevalent in exhibition feedback given the diversity of our participant sample. Unlike more complex models such as BERT (*59*), VADER uses a pre-trained lexicon instead of requiring extensive training on large labeled datasets, making it an efficient and practical choice for our analysis.

Case study 2: Reflections on user agency and interaction modalities

One interesting question is how much agency users have in shaping the swarm behavior. In this work, we investigate three modes of operation: passive observation, crowd interaction, and creative expression, with user agency increasing from less to more, respectively. In the passive observation mode, users have no control and simply watch the system's pre-programmed behaviors. In the crowd interaction mode, users can influence individual robots, with the swarm's emergent behavior reflecting the collective input of all participants. In the creative expression mode (using the wearable), users have the highest level of agency, creating large-scale coordinated effects by influencing multiple robots simultaneously.

To give users even more agency and make the Swarm Garden more intuitive and inclusive, we are exploring a range of future interaction modalities that aim to provide natural, embodied, and versatile ways for users to interact with the swarm. Specifically, we present three modalities: wearable, speech, and tablet interfaces.

Wearable: The wearable provides a platform to explore controlling the swarm in an embodied and accessible way. It can be further adapted to support different accessibility needs, such as controlling the Swarm Garden with different body parts or different movements.

Speech: We are exploring speech-based interaction using Large Language Models to interpret verbal commands. This allows users to intuitively control system behaviors such as adjusting shading levels, setting the mood through LED colors, or initiating interaction modes, without the need for physical interfaces. This modality would particularly benefit users with visual impairments, as it removes reliance on screens or visual cues.

Tablet interface: To support users with limited mobility, a tablet-based Graphical User Interface (GUI) could simulate the Swarm Garden. This would allow users to select and control specific SGbots individually or in groups, providing an alternative to gesture-based interaction.

Swarm Garden provides an exciting avenue for exploring user agency in adaptive façades. Most existing reactive installations are rigid: either fixed physically or limited to pre-programmed interactions, giving users little to no agency in shaping system behavior. Swarm Garden is fully programmable, capable of detecting a wide range of stimuli (for example, gestures, sounds, colors, faces, or signals from wearables) and responding with customized LED and blooming patterns. This

programmability opens opportunities for interactions we are yet to explore as well as adaptation over time by learning from individual or collective behaviors to enhance engagement. Users can also physically reconfigure the system, changing its shape or relocating it with ease thanks to its modular and portable design. Unlike light-based reactive installations that users can affect only in limited ways, our system gives both us and the users the flexibility to rearrange the robots, switch between behaviors, connect to external devices, and enable or disable interaction, allowing the swarm to change its appearance and behavior from day to day.

New work further explores how the system can learn from users over time, enabling them to train and adapt the system to their preferences. This could lead to further user agency in controlling long-term behaviors of the system.

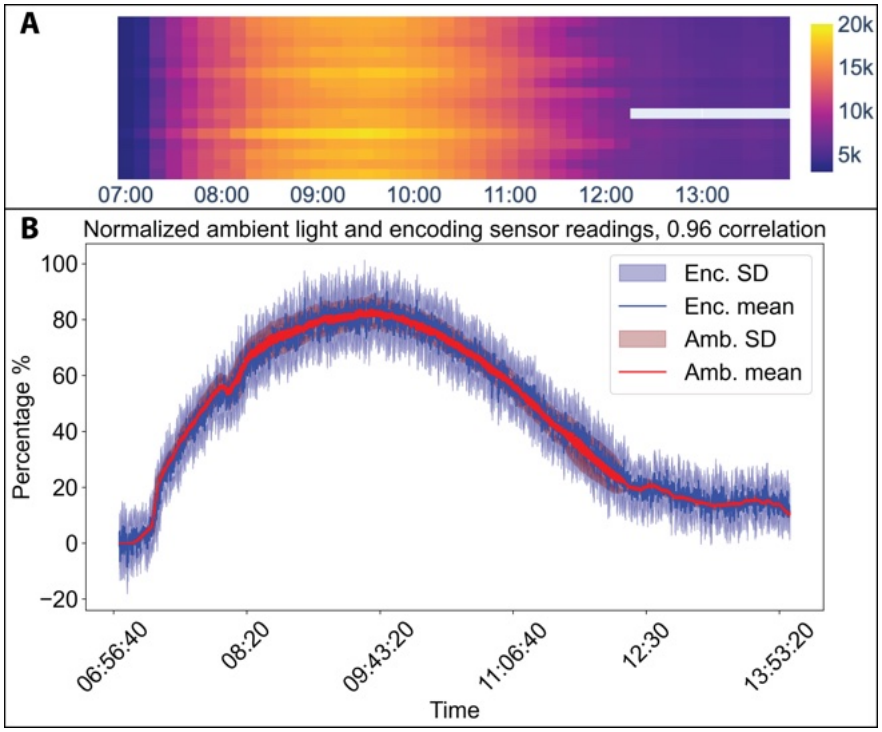

Fig. S1. **Third day's experiment results for the adaptive shading individualistic approach.** (**A**) Heatmap showing the raw illuminance values measured by each SGbot during the experiment. Each row in the heat map corresponds to 1 of 16 SGbots. (**B**) The 10-min running means of the normalized encoding (Enc.) sensor readings, corresponding to plate displacement, are shown as a solid blue line averaged across all 16 robots, with the shaded blue region indicating the SD from the mean across all robots (n = 16). The 10-min running means of the normalized ambient (Amb.) light sensor readings are shown as a solid red line averaged across all 16 robots, with the shaded red region indicating the SD from the mean across all robots (n = 16). The empty values in the heatmap (in gray) indicate loss of communication with one of the robots for the last 2 hours of the experiment.

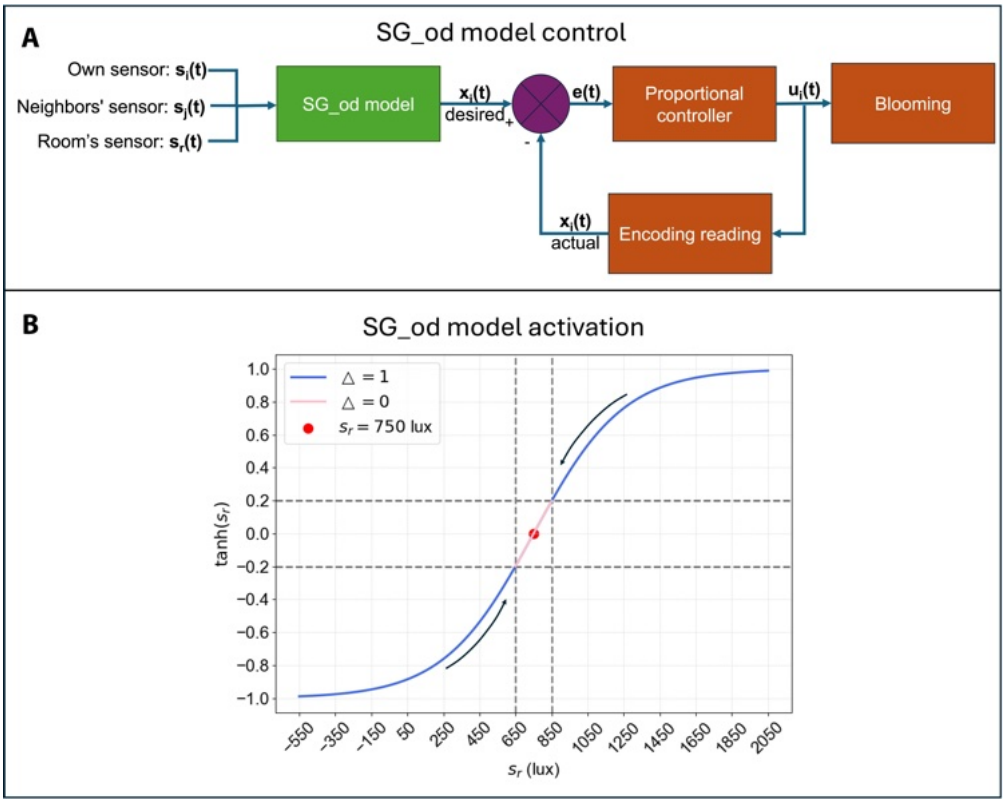

Fig. S2. **Collective adaptive shading control and SG_od model activation.** (**A**) SG_od's control system schematic. The system receives information from three sources of input: its own ambient light sensor, the ambient light sensors of neighboring robots, and the room's ambient light sensor. Based on these inputs, the SG_od model computes the desired blooming level, which is then passed to a proportional controller to achieve it. (**B**) The hyperbolic tangent function (scaled and with normalized room's ambient light sensor's input) that determines whether the SG_od model is deactivated (pink line) or activated (blue line). The more the room's illuminance deviates from 750 lux in either direction, denoted by the black arrows, the greater the contribution of the third term in Eq. 2.

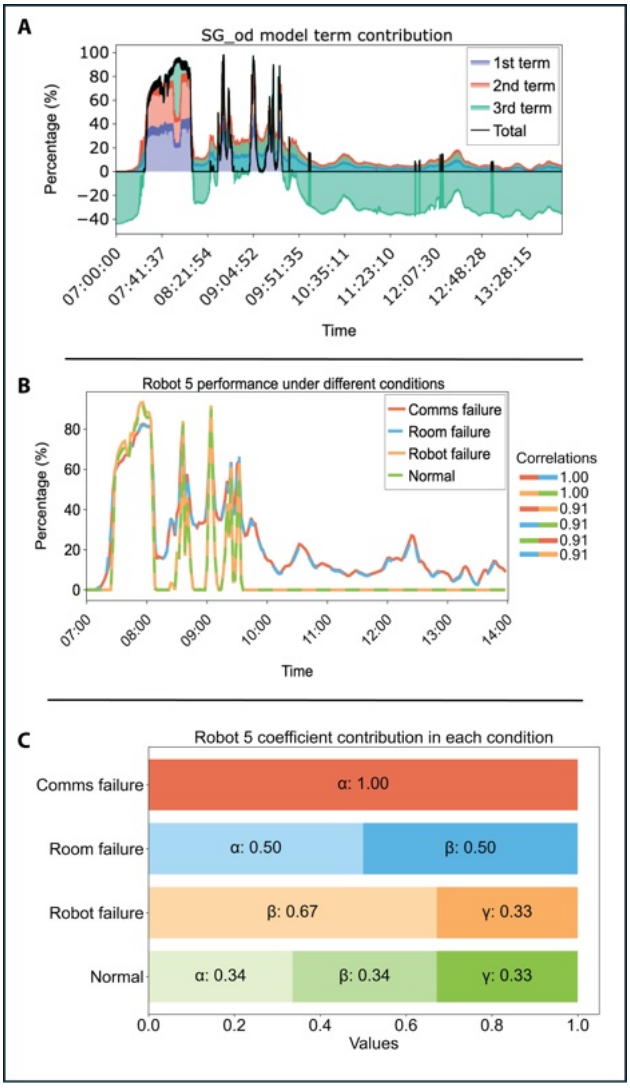

Fig. S3. **Third day's adaptive shading experiment results (rainy day) using collective opinion dynamics run from 7:00 to 14:00 in a real office.** (**A**) Each term's contribution from the SG_od model (see Eq. 2) represented using a different color, with the final value (total) $x_i(t)$ represented in black. (**B**) Final value, $x_i(t)$, when introducing four failure scenarios in simulation to SGbot 5 for the same day: Red is for communication failure, blue is for the room's ambient light sensor failure, orange is for SGbot's ambient light sensor failure and green is for no failures (normal scenario). The Pearson correlation coefficients between each of the four scenarios are shown on the right. (**C**) Average value of each coefficient contribution (used in Eq. 3) in each failure scenario.

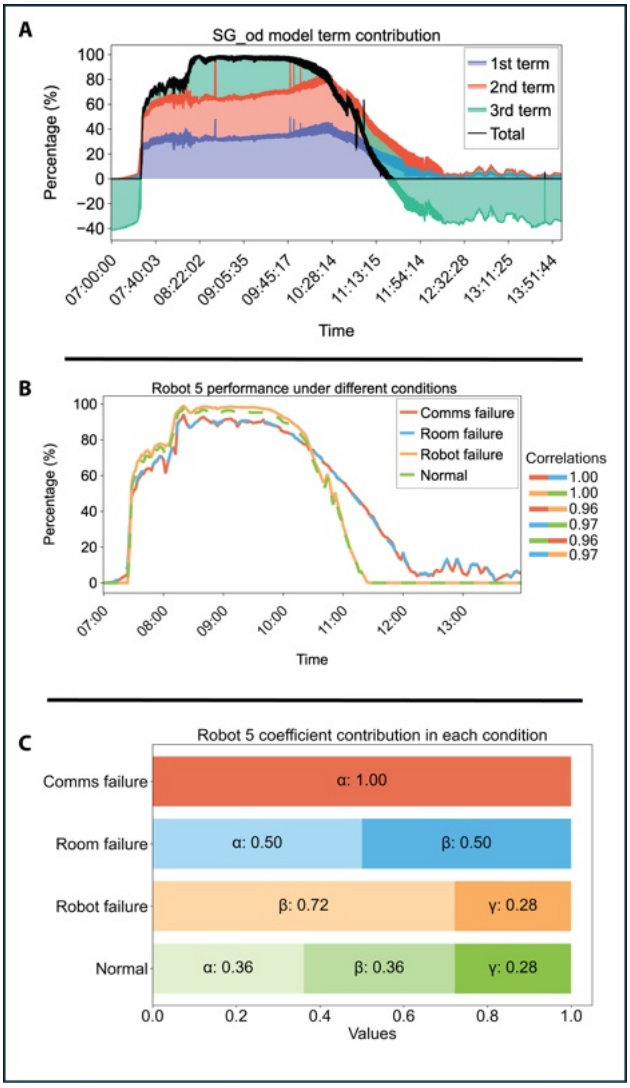

Fig. S4. **Fourth day's adaptive shading experiment results (sunny day) using collective opinion dynamics run from 7:00 to 14:00 in a real office.** (**A**) Each term's contribution from the SG_od model (see Eq. 2) represented using a different color, with the final value (total) $x_i(t)$ represented in black. (**B**) Final value, $x_i(t)$, when introducing four failure scenarios in simulation to SGbot 5 for the same day: Red is for communication failure, blue is for the room's ambient light sensor failure, orange is for SGbot's ambient light sensor failure and green is for no failures (normal scenario). The Pearson correlation coefficients between each of the four scenarios are shown on the right. (**C**) Average value of each coefficient contribution (used in Eq. 3) in each failure scenario.

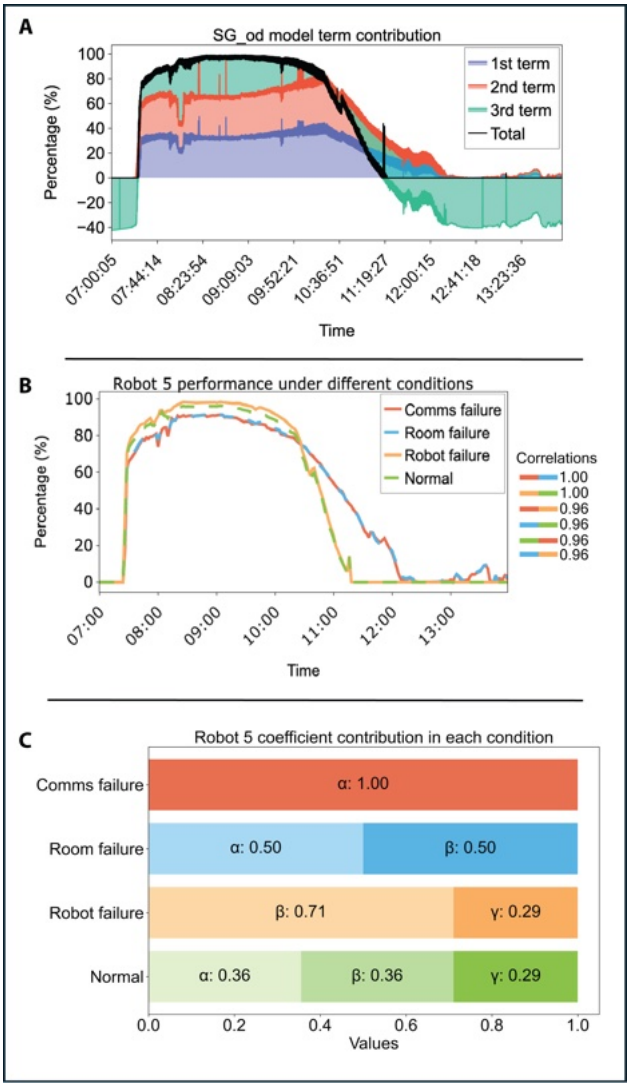

Fig. S5. **Fifth day's adaptive shading experiment results (sunny day) using collective opinion dynamics run from 7:00 to 14:00 in a real office.** (**A**) Each term's contribution from the SG_od model (see Eq. 2) represented using a different color, with the final value (total) $x_i(t)$ represented in black. (**B**) Final value, $x_i(t)$, when introducing four failure scenarios in simulation to SGbot 5 for the same day: Red is for communication failure, blue is for the room's ambient light sensor failure, orange is for SGbot's ambient light sensor failure and green is for no failures (normal scenario). The Pearson correlation coefficients between each of the four scenarios are shown on the right. (**C**) Average value of each coefficient contribution (used in Eq. 3) in each failure scenario.

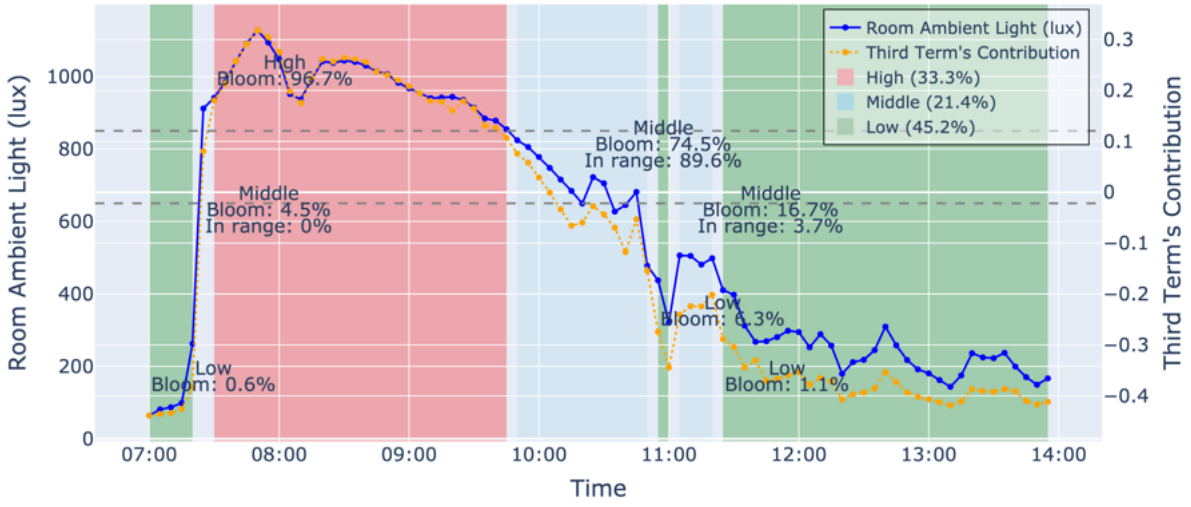

Fig. S6. **First day's room ambient light results (sunny day) using collective opinion dynamics.** The blue line represents the room ambient light level in lux (the raw sensor values measured by the room sensor) from 7:00 to 14:00 (left y axis). The yellow line represents the third term's contribution (in Eq. 2) in the SG_od model (right y axis), which were values collected throughout the experiment. The High region is shown in red (occurring 33.3% of the time), the Low region in green (45.2% of the time), and the Middle region in light blue (21.4% of the time). On the figure, the average blooming level (how buckled or flattened the robots were) is shown in black text for each period of High, Low, and Middle. Additionally, for the Middle periods, the percentage of time when the ambient light stayed within the preferred range of 650–850 lux is shown.

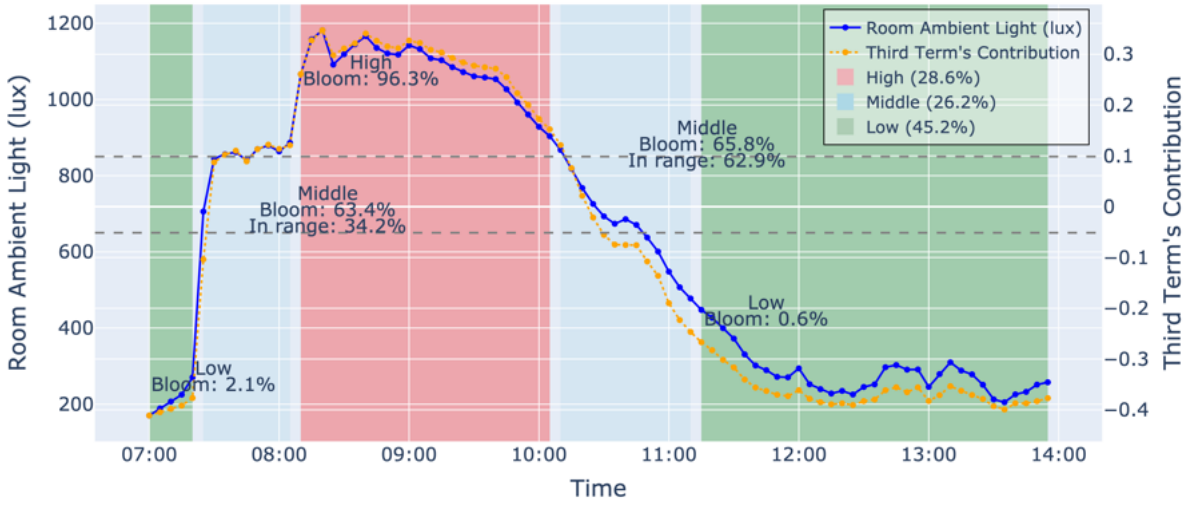

Fig. S7. **Fourth day's room ambient light results (sunny day) using collective opinion dynamics.** The blue line represents the room ambient light level in lux (the raw sensor values measured by the room sensor) from 7:00 to 14:00 (left y axis). The yellow line represents the third term's contribution (in Eq. 2) in the SG_od model (right y axis), which were values collected throughout the experiment. The High region is shown in red (occurring 28.6% of the time), the Low region in green (45.2% of the time), and the Middle region in light blue (26.2% of the time). On the figure, the average blooming level (how buckled or flattened the robots were) is shown in black text for each period of High, Low, and Middle. Additionally, for the Middle periods, the percentage of time when the ambient light stayed within the preferred range of 650–850 lux is shown.

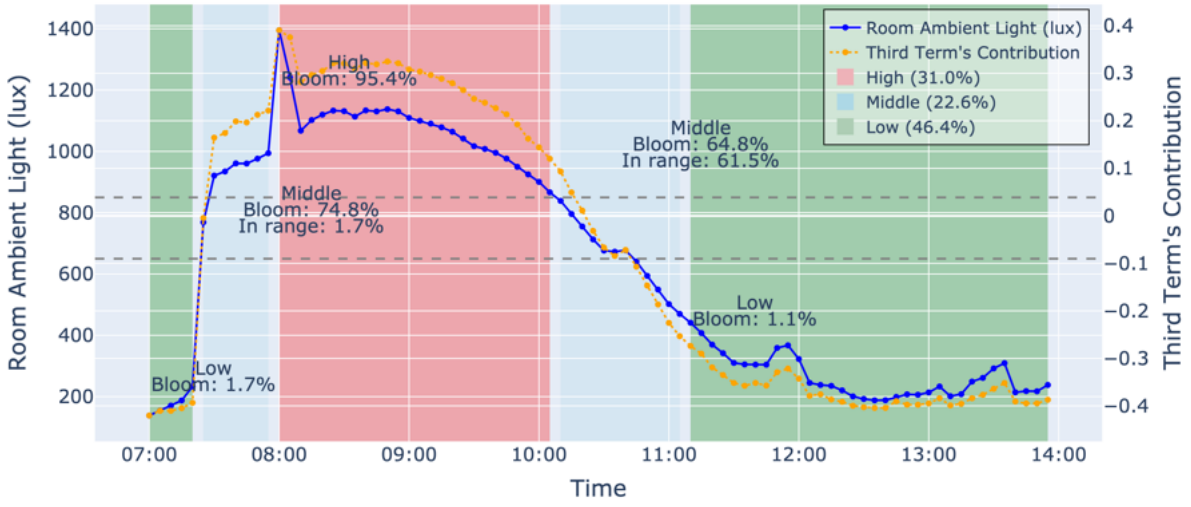

Fig. S8. **Fifth day's room ambient light results (sunny day) using collective opinion dynamics.** The blue line represents the room ambient light level in lux (the raw sensor values measured by the room sensor) from 7:00 to 14:00 (left y axis). The yellow line represents the third term's contribution (in Eq. 2) in the SG_od model (right y axis), which were values collected throughout the experiment. The High region is shown in red (occurring 31% of the time), the Low region in green (46.4% of the time), and the Middle region in light blue (22.6% of the time). On the figure, the average blooming level (how buckled or flattened the robots were) is shown in black text for each period of High, Low, and Middle. Additionally, for the Middle periods, the percentage of time when the ambient light stayed within the preferred range of 650–850 lux is shown.

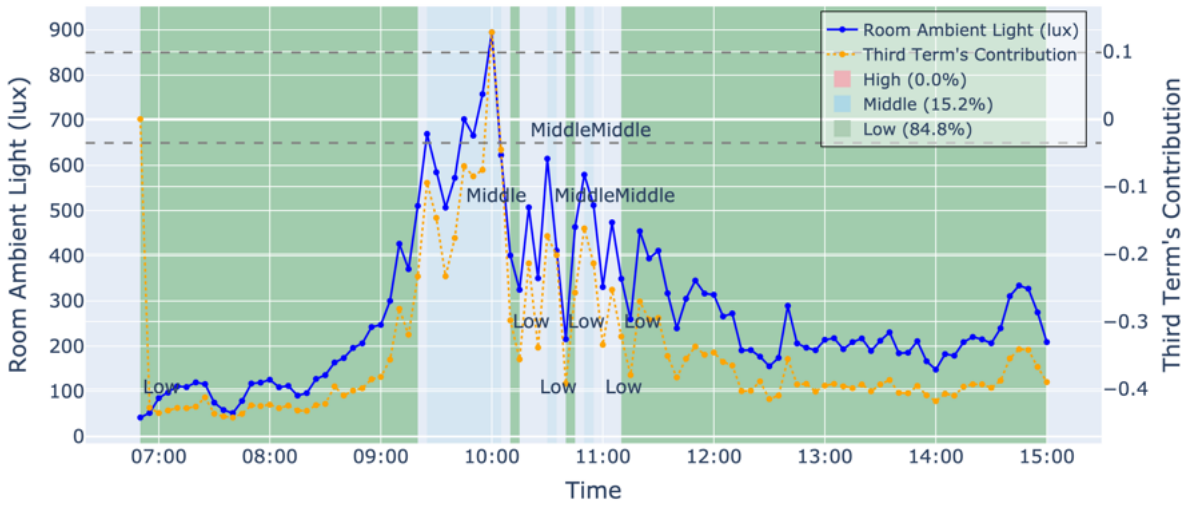

Fig. S9. **Second day's room ambient light results (cloudy day) using collective opinion dynamics.** The blue line represents the room ambient light level in lux (the raw sensor values measured by the room sensor) from 7:00 to 14:00 (left y axis). The yellow line represents the third term's contribution (in Eq. 2) in the SG_od model (right y axis), which were values collected throughout the experiment. The system never experienced the High period due to the weather conditions which otherwise would be shown in red (occurring 0% of the time), the Low region is shown in green (84.8% of the time), and the Middle region in light blue (15.2% of the time). The bloom and preferred range percentages are omitted to avoid text clutter.

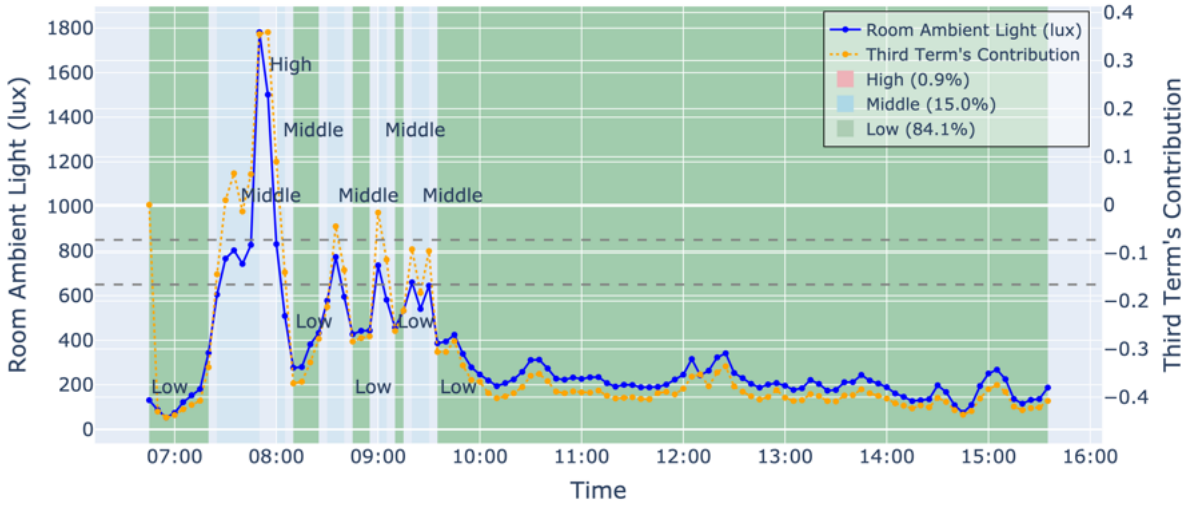

Fig. S10. **Third day's room ambient light results (rainy day) using collective opinion dynamics.** The blue line represents the room ambient light level in lux (the raw sensor values measured by the room sensor) from 7:00 to 14:00 (left y axis). The yellow line represents the third term's contribution (in Eq. 2) in the SG_od model (right y axis), which were values collected throughout the experiment. There was a very brief moment of a High period, which was not long enough to be shown on the graph (occurring 0.9% of the time), the Low region in green (84.1% of the time), and the Middle region in light blue (15% of the time). The bloom and preferred range percentages are omitted to avoid text clutter.

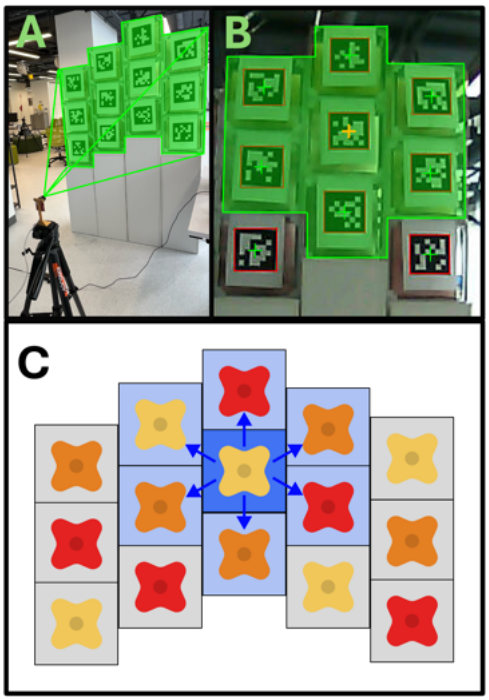

Fig. S11. **Localization and neighbor identification setup.** (**A**) Camera capturing SGbots' IDs through their AprilTags. (**B**),(**C**) Example of an SGbot's neighborhood (shown for the middle SGbot).

Table S1: **SGbot's bill of materials.** In future versions, the Nicla Vision can be replaced with a cheaper alternative, such as an ESP32, combined with only the sensors necessary for the intended application. This would reduce cost and make the system more scalable.

Category	**Component**	**Price**	**Subtotal**
Mechanical	Stepper motor	$11.99	$39.26
	Shaft coupling	$1.50	
	Threaded rod	$9.80	
	Threaded rod adhesive nut	$2.17	
	Screws	$0.60	
	Plastic shimstock sheets	$13.20	
Electrical	Ribbon cable	$0.86	$34.61
	IDC sockets	$1.81	
	MicroUSB socket	$0.28	
	Terminal blocks	$0.23	
	Switches	$1.00	
	JST-XH connectors and cable connectors	$1.10	
	Power bank	$22.00	
	LED Strip	$4.00	
	L-shaped LED strip connectors	$3.33	
Control & Sensors	Arduino Nicla Vision	$115.00	$131.65
	I/O expander	$5.95	
	Motor driver	$2.87	
	Time-of-flight sensor	$2.83	
	Ambient light sensor	$5.00	
Structural	Acrylic sheets	$26.66	$107.54
	Miniature t-slots aluminum extrusions	$49.48	
	Corner t-slot screws	$31.20	
	Nuts for miniature t-slots	$0.20	
		Grand Total	**$313.06**

Table S2: **Swarm Garden control latency.**

Component	**Control latency (s)**	**Comments**
Interaction sensor	$\mu = 0.35$, $\sigma = 0$	5 robots × 30 samples
Ambient light sensor	$\mu = 0.03$, $\sigma = 0.02$	5 robots × 30 samples
Wearable	Z-axis Down: $\mu = 1$, $\sigma = 0$ Z-axis Up: $\mu = 1$, $\sigma = 0$ X-axis Up: $\mu = 1$, $\sigma = 0$ Y-axis Down: $\mu = 1.03$, $\sigma = 0.01$ Y-axis Up: $\mu = 1.02$, $\sigma = 0.01$	20 samples per axis. A 1s delay was added to enable the swarm to execute the response (such as light propagation) to ensure observers could clearly see each response without overlap.
Blooming	Actuate for 1 second at a time before checking encoding sensor	To maintain effective closed-loop control, the actuator moves in 200 step increments at 5ms per step (1s total) before checking the encoding sensor. This intentional rate limits jerky motion between sensor readings for a smooth and effective closed-loop control.
LED	0.01-1.275. Swarm-wide color propagation (rightmost to leftmost robot across 12 columns): 0.12-15.3	Adding brief delays ensures smooth, gradual color transitions rather than abrupt changes that are visually jarring. The fade duration depends on the distance between the current and target colors: it can be as short as 0.01s or as long as 1.275s per robot.
Messaging	Sending: $\mu = 0.0000826$, $\sigma = 0.0000169$. Receiving: $\mu = 0.04734$, $\sigma = 0.04763$	2 robots × 30 samples sending and acknowledging messages. In practice, some delays were added - for example, messages were sent only after LED execution to maintain a pleasing and human-perceivable rhythm (exhibition), or every 30s (shading) to prevent network overload.

Table S3: **Battery life for different modes.**

Mode	**Component**	**Experiment**	**Average**
Blooming	Motor + encoding sensor	Three robots continuously blooming for one hour.	μ = 9%/hour, σ = 0. Linear approximation: 11 h 6 min to full depletion.
LED	LED + interaction sensor	Three robots continuously fading between different colors and checking interaction sensor.	μ = 14%/hour, σ = 1.41% Linear approximation: 7 h 8 min to full depletion.
Shading	Motor when adjustment was needed, encoding sensor, ambient light sensor, networking, opinion dynamics model, on-board Nicla LED.	Data collected over five days of collective dynamics deployments with 16 robots. Each robot began at 100% charge, experiments continued until depletion; only 7 hours of operation were included in analysis.	μ = 7 h 47 min, σ = 40 min to full depletion. The wireless system and on-board Nicla LED also consumed additional energy; battery depletion is likely non-linear.

Table S4: **High-level metric comparison between individualistic and collective approaches.** Collective strategies generally outperform individualistic approaches in fault tolerance, robustness, and user preference incorporation, while both respond effectively to environmental changes.

Metric	**Individualistic**	**Collective**
Responsiveness to environment	Yes	Yes
Fault tolerance	No	Yes
User preference	No	Yes
Robustness	No	Yes

Table S5: **Atrium configurations' contrast ratio.**

Description	**Min lux**	**Max lux**	**Ratio**
No Modules (Fig. 5Ci)	3235.99	7514.73	2.3
Closed (Fig. 5Cii)	439.79	1313.35	3.0
Half (Fig. 5Ciii)	1318.27	3220.37	2.4
Open (Fig. 5Civ)	2288.72	5423.32	2.4
Horizontal (Fig. 5Cv)	1109.37	3508.78	3.2
Vertical (Fig. 5Cvi)	1097.56	3470.83	3.2
9h (Fig. 5Di)	851.30	2557.83	3.0
13:00 (Fig. 5Dii)	776.55	2625.10	3.4
17h (Fig. 5Diii)	909.77	2288.52	2.5

Table S6: **Comparisons of the Atrium experiment against static configurations.**

Time	**Static**	**Adaptive**	**Avg. diff.**
9:00	No modules	Fig. 5Di	120.7%
13:00	No modules	Fig. 5Dii	108.2%
17:00	No modules	Fig. 5Diii	52.8%
9:00	Open modules	Fig. 5Di	96%
13:00	Open modules	Fig. 5Dii	81.1%
17:00	Open modules	Fig. 5Diii	18.7%

Caption for Movie S1. Adaptive shading individualistic approach experiment on a sunny day. Time-lapse showing the 7-hour experiment (compressed into 30 seconds) presented in Fig. 3A.

Caption for Movie S2. Adaptive shading individualistic approach experiment on a cloudy day. Time-lapse showing the 7-hour experiment (compressed into 30 seconds) presented in Fig. 3B.

Caption for Movie S3. Adaptive shading collective opinion dynamics experiment on a sunny day. Time-lapse showing the 7-hour experiment (compressed into 30 seconds) presented in Fig. 4A.

Caption for Movie S4. Adaptive shading collective opinion dynamics experiment on a cloudy day. Time-lapse showing the 7-hour experiment (compressed into 30 seconds) presented in Fig. 4B.

Caption for Movie S5. Passive observation patterns. An example run of the passive observation patterns, shown in Fig. 6C. SGbots were not yet equipped with the railing system that prevents the plate from rotating off-axis.

Caption for Movie S6. Selective interaction mode used in crowd interactions. A demonstration showing the interaction mode used in the public exhibition (see Fig. 6D-F) that allows users to bloom the robots with hand gestures.

Caption for Movie S7. Additional selective interaction mode: Transient light trail. A demonstration showing an additional selective mode: as the operator moves along the wall, a transient trail of light follows their path, creating a dynamic, shadow-like effect.

Caption for Movie S8. Additional selective interaction mode: Light painting. A demonstration showing an additional selective mode: using the wearable device, a user can change individual SGbot LED colors to customize their interior space.

Caption for Movie S9. Improvisational dance performance. A performance created by A.J. equipped with the wearable device in collaboration with the Swarm Garden.

Caption for Movie S10. The wearable device interactions. Responses of the Swarm Garden from certain interactions involving arm movement caused by users equipped with the wearable device.

Caption for Data S1. Exhibition word cloud responses. This file contains all the responses received for the word cloud from attendees. Each attendee was allowed up to three words.

Caption for Data S2. Exhibition open feedback responses. This file contains all the feedback responses received from attendees. Personal or identifying remarks are redacted for confidentiality.