

PRINTING LIVING WALLS

ADDITIVE MANUFACTURING OF POROUS
CERAMICS FOR A SELF-SUSTAINING
BIORECEPTIVE FACADE



ADDITIVE MANUFACTURING OF POROUS CERAMICS FOR A SELF-SUSTAINING BIORECEPTIVE FACADE

Abstract

Cities absorb and retain solar heat, intensifying the urban heat island effect at a growing cost to health and liveability. Vegetated surfaces cool their surroundings through evaporation, and mosses are well suited to colonising vertical surfaces, yet as poikilohydric organisms they cannot regulate or store their own water. During the summer heatwaves when cooling is most needed, an unsupplied moss layer dries out, falls dormant, and can leave a surface hotter than bare concrete.

This thesis develops a bioreceptive, 3D-printed porous ceramic facade tile that acts as an artificial vascular system, harvesting and storing rainwater and feeding it back to the biological layer by capillary action so the moss stays active through dry periods, while also providing a substrate for urban biodiversity. The work follows a Research-through-Design methodology across three scales. At the micro-scale, a white stoneware body was engineered with spent coffee grounds and corn starch as sacrificial pore-formers and fired at 1100 °C, producing a pH-neutral ceramic with an apparent porosity of about 39 % and a bimodal pore structure predicted to resist frost.

At the meso-scale, five printed surface geometries were compared for capillary uptake, moisture retention and rainfall capture, and a fractal-branching geometry showed the most promising combination, though single-sample testing makes the comparison indicative. These geometries were produced through a parametric pipeline in which a Python script generates the print toolpath directly and an interactive viewer lets each tile be previewed and exported for fabrication.

At the macro-scale, the validated tile was integrated, as a conceptual application, into a rail-mounted, gravity-fed module with an internal reservoir and a partial glaze that concentrates moisture in the porous zones where moss establishes. Outdoor testing showed the wet ceramic running far cooler than a conventional wall and staying close to air temperature through the day, which locates the cooling in the water-holding body rather than the plant; the glaze and the moss each give up a little cooling for a large gain in water retention. Transplanted moss established and has stayed alive for one month and seven days at the time of writing, healthy and still going, and a chamber test showed a single reservoir fill sustaining the module for roughly a day under worst-case drying, pointing to a larger store or passive rainwater buffering as the decisive step toward a genuinely self-sustaining facade.

Acknowledgements

This project grew out of the things I most wanted to spend 20 weeks on: working with my hands in ceramics and additive manufacturing, building the computational tools that shape them, and doing it in service of a problem that matters, the heat our cities now hold. What began as a curiosity about whether a printed surface could keep moss alive became an attempt to carry one idea all the way from a raw clay recipe to a living prototype. This was not something I did alone. The people here are the reason it became what it is.

First, I want to thank my supervisors, Dr. S. Parisi and Dr. M. Özdemir, for their guidance, their critical questions, and their encouragement throughout the project. Their feedback shaped both the direction of the research and the way it is presented here.

I am grateful to Max Latour of Urban Reef, who welcomed me into the workspace where he makes his pieces for Urban Reef and openly shared how he works, including his print settings and approach, while answering my many questions along the way. His openness and experience with bioreceptive design were a real help in getting the project off the ground.

My thanks go to Vera and Henri at the LAMA labs, who helped me with the printing and were proactive, patient, and genuinely kind throughout. Working alongside them made the most demanding part of the fabrication far easier than it would otherwise have been.


I would also like to thank Tim de Jong at the Applied Labs, who helped me think through how to run several of the tests and patiently explained the equipment and methods I was new to.

I am thankful to my parents, for supporting me financially with materials and supplies, and for their help transporting the prototypes when I could not manage it on my own.

Finally, I want to thank my girlfriend, Jasmijn Jiskoot, for coming along to collect moss, helping take the photographs, and assisting wherever she was needed. Her patience and support meant a great deal over the course of this work.

Contents

01	INTRODUCTION 1.1 Global Context 1.2 Material Bioreceptivity & Bryophyte Ecology 1.3 Unplanned Bioreceptivity — Field Research 1.4 State-of-the-Art Review	06
02	AIM & RESEARCH PLAN 2.1 Project Objective 2.2 Research Framework	18
03	PHASE 1 — MICRO-SCALE 3.1 Introduction 3.2 Literature Review 3.3 Research Questions 3.4 Methodology 3.5 Results & Discussion	22
04	PHASE 2 — MESO-SCALE 4.1 Introduction 4.2 Background 4.3 Geometric Design & Fabrication 4.4 Methodology 4.5 Results & Discussion	40
05	PHASE 3 — MACRO-SCALE 5.1 Introduction 5.2 Systemic Integration 5.3 Methodology 5.4 Results & Discussion	54
06	CLOSING 6.1 Conclusion 6.2 Reflection	76
07	REFERENCES	84
08	APPENDICES A Sample catalogue B Geometry & water behaviour C Module longevity & pH D Parametric pipeline code E Thermal measurements	86 86 90 92 94 98





01

INTRODUCTION

Urban Heat Island · Bioreceptivity · Project Objective

Fig. 1.1 — Moss growth between a rusted steel plate.

This section sets out why urban heat is a problem worth designing against, why vegetated surfaces are a credible part of the response, and why the building facade in particular is where that response has room to grow. It closes on the limitation that defines the rest of the thesis: this cooling only works while the plant is alive and evaporating.

1.1

Global Context

Modern cities are, in effect, engineered to trap heat. Concrete structures and asphalt surfaces absorb solar radiation through the day and store it like a thermal battery, then release it slowly after sundown, while densely packed blocks keep the air stagnant and warm into the night. The result is the Urban Heat Island Effect, in which urban areas run consistently hotter than the rural land around them because their roads, roofs and walls retain heat that vegetation would otherwise moderate (Asian Development Bank, 2024). The effect tends to compound itself, since rising temperatures push up demand for air conditioning, which raises energy use and the emissions that drive further warming (Nuruzzaman, 2015). The human cost is not trivial. Heat is now a leading climate-related cause of death, and the burden falls hardest on low-income neighbourhoods, which typically hold more heat-retaining material and less greenery than wealthier districts (Zheng et al., 2026).

In this context, vegetation has shifted from an amenity to a recognised instrument of climate adaptation. Plants cool their surroundings through evaporation, a process in which liquid water draws in heat as it shifts to vapour and carries that energy away from the surface, producing a measurable local drop in temperature (Sustainability Directory, 2026). What was once treated as a pleasant side effect of greenery is increasingly quantified as infrastructure. Modelling that links neighbourhood greening to community health outcomes finds that even modest increases in natural features cut heat exposure during heatwaves, reducing the daily humidex by more than eleven percent and lowering the associated health risks that follow extreme heat (Dardir, Wilson, & Berardi, 2026). Policy has begun to follow the evidence, with the UK Green Building Council calling for nature-based solutions across all new buildings and infrastructure by 2030 so that cities stay habitable as they warm (Ing & Jessel, n.d.).




Fig. 1.2 — Urban heat island in delft, IR photograph.

The cooling itself is well documented at the scale of the building envelope. Reviews of exterior vertical greening report reductions in surrounding air temperature of up to roughly 3 °C close to the wall, with the effect measurable up to about 0.6 metres away from the surface (Green Facades and Living Walls, 2019). At the scale of the whole city, simulations suggest that widespread vertical greening can lower overall urban air temperature by around 0.7 to 0.9 °C, enough to cut the intensity of the heat island effect by almost half (Comprehensive Assessment of the Impact of Green Roofs and Walls, 2024). The principle holds for moss as well as for vascular plants: paired measurements on concrete cubes found that a moss-covered face sat at 28.6 °C against 30.5 °C for plain concrete outdoors, and lowered the temperature behind the wall from 28.0 °C to 27.0 °C indoors (Pathak & Brahmandam, n.d.).

While a single facade cannot solve the urban heat crisis on its own, the data demonstrates that living surfaces are a proven, effective tool for localized temperature reduction. This project focuses on delivering realistic, incremental cooling rather than presenting a single cure-all.

If vegetated surfaces are part of the answer, the question becomes where to put them. Dense urban areas have almost no spare ground for planting, yet they are surrounded by vertical surface that does nothing thermally at all. Facades are a vast and largely untapped spatial resource, and turning these otherwise sterile concrete and brick walls into living surfaces captures urban cooling without competing for scarce real estate at street level. This reframes the facade from a passive boundary into a working part of the city's climate response.

That promise comes with a condition, and the condition is the starting point for everything that follows. Evaporative cooling only happens while the plant is actively transpiring; the moment it dries out and goes dormant, the cooling stops, and as later sections will show, a dry biological layer can make the surface worse rather than better. The central problem of this thesis is therefore not whether a living facade can cool, but how to keep it alive and evaporating through exactly the dry, hot periods when cooling is needed most. The next section examines the organism that makes this possible, the property that makes a surface colonisable, and the physiological limit that the rest of the project is built to overcome.

1.2

Material Bioreceptivity & Bryophyte Ecology




Fig. 1.3 — Close-up macro photography of moss leaves.

This section introduces the organism at the centre of the project and the material property that lets it grow on a wall, then sets out the physiological limit that the rest of the thesis is built to overcome. It ends by stating the core design problem and the two contributions the system is intended to make.

For a wall to host life at all, its surface has to be hospitable, and the degree to which it is can be treated as a property of the material rather than an accident. René Guillitte named this property bioreceptivity in 1995, defining it as the measurable capacity of an inert material to be colonised by living organisms without being structurally degraded in the process. The concept marks a shift in how biological growth on buildings is understood. What was traditionally classed as biodeterioration, a defect to be cleaned off and prevented, is reframed as an ecological asset that a surface can be deliberately designed to invite (Pollini & Rognoli, 2021). Designing for it means working backwards from the needs of the organism rather than treating colonisation as something that simply happens to a material.

The organisms best suited to a vertical surface are bryophytes, the mosses, together with lichens. Their advantage begins with how they attach. Unlike vascular plants, mosses are rootless and anchor themselves with rhizoids, fine root-like structures that grip the surface purely for mechanical hold and draw neither water nor nutrients from it (Veeger et al., 2023). Because nothing penetrates the substrate in search of

moisture, growth does not drive the cracking and material breakdown normally associated with biological colonisation, which is what makes moss compatible with a building envelope rather than a threat to it (Cruz & Beckett, 2016).

That same rootless construction is also the source of the project's central difficulty. Mosses are poikilohydric, meaning they cannot regulate their own internal water content and their hydration tracks the surrounding environment directly (Jakubovskis, 2023). Vascular plants control water loss through stomata, microscopic valves that throttle evaporation to balance hydration; mosses have no such mechanism and instead evaporate passively across their whole surface, their leaves often only a single cell thick (Proctor, 1982; Veeger et al., 2023). This cuts two ways. On one hand it makes moss an unusually efficient evaporative surface, since the entire organism is active cooling area rather than a regulated leaf canopy. On the other it leaves the plant defenceless against drought. During a summer heatwave, exactly when cooling is most needed, an unprotected moss layer

quickly exhausts its moisture, enters metabolic dormancy, and stops evaporating altogether (Veeger et al., 2023).

The consequence of that dormancy is what turns the problem from a limitation into a liability. Veeger et al. (2026) found that a dry, dormant moss layer under direct sun does not merely stop cooling but actively heats the surface, running between 1.5 °C and 4.1 °C hotter than bare concrete during the day. At night the same layer works against the building a second time, its low thermal conductivity insulating the wall and trapping stored heat so that the substrate behind it sits 3.0 °C to 3.8 °C warmer than an uncovered equivalent. A moss facade that is allowed to dry out therefore performs worse than no facade at all. Sustained moisture is not a refinement of the concept; it is the condition on which the entire thermal benefit depends.

This is the core challenge the thesis addresses. Because the organism can neither store water nor reach the groundwater a vascular plant would tap, the burden of water management falls entirely on the surface it lives on.




Fig. 1.4 — Close-up macro photography of moss leaves.

The material itself must act as an artificial vascular system: harvesting rainwater, holding it in reserve, and feeding it back to the biological layer by capillary action so the moss stays active through the dry intervals between rain. Delivering that function through the material and geometric freedom of ceramic additive manufacturing is the proposition the rest of this work develops.

Bioreceptivity, it should be noted, is governed by more than water. Surface chemistry and pH, surface roughness and pore geometry for anchorage, shelter from UV and wind, and the capture of airborne nutrients all shape whether an organism establishes (Parveen & Tariq, 2023). This thesis treats several of these as properties to be satisfied rather than varied: a biologically safe pH and an appropriately rough, porous surface are designed into the material in Phase One and then held constant. The variable the project isolates and builds its structure around is water, because the field evidence and the existing engineered systems reviewed in the next sections both point to sustained moisture as the factor that actually decides whether a living facade survives the conditions it was built for.

Keeping moss alive yields a second benefit beyond cooling, and it is treated here as a goal in its own right rather than a side effect. A hydrated moss layer is a foundation for urban biodiversity. Mosses have been found to harbour almost twice the microbial richness and diversity of vascular plants, precisely because their lack of vasculature and protective cuticle leaves their tissues open to colonisation by a far broader environmental pool, while dense moss mats trap thin films of water in which bacterial communities establish (Jakubovskis, 2024). Over time these communities, together with the dust and organic debris the surface accumulates, build new ecological niches that come to support invertebrates such as mites, springtails and small insects, which in turn anchor more complex interactions (Parveen & Tariq, 2023). A living facade in this sense does two things at once: it cools the surface it covers, and it converts a sterile wall into habitat in cities that have very little of it left.

The question that follows is how to keep such a surface reliably wet. That answer does not have to be invented from scratch, because moss already colonises hard urban surfaces on its own wherever the conditions happen to be right. The next section examines where and why this spontaneous colonisation occurs, and distils it into a set of design principles for engineering those conditions deliberately.




Fig. 1.5 — Close-up macro photography of moss leaves.

1.3

Unplanned Bioreceptivity

Field Research

The goal of the field research was to analyse where and why moss colonises urban surfaces on its own, and to translate those growth patterns into criteria for designing a bioreceptive surface. Throughout the city, moss already establishes on untreated concrete, brick and metal that was never intended to support life, and the distribution is far from random. A local survey documented this unplanned colonisation, recorded in the photographs opposite, reading each surface for what allowed growth or prevented it. The patterns point to one consistent conclusion. Colonisation is governed less by the material than by the microclimate each surface creates. Rain sheets straight off smooth glass and metal and nothing holds, whereas moss concentrates wherever the geometry slows water down, shelters it from sun and wind, and traps the fine debris that becomes a rooting medium. A single surface can be bare or thick with growth along its length depending only on where water is made to linger. This is what the field study contributes to the design work: the task is not to find a special material but to shape a surface that lengthens the time water stays, shades it once it lands, and holds the sediment an organism needs to root.



OBS. 03 — LUSH GROWTH ON CURB



OBS. 06 — ECOLOGICAL SUCCESSION GRADIENT




OBS. 05 — GROWTH BETWEEN THE CRACKS

Fig. 1.6 — Photo collage of unplanned bioreceptivity in the urban environment.



OBS. 01 — BIOFILM ON SMOOTH SURFACE



OBS. 02 — CORRUGATED GEOMETRY



OBS. 04 — POROUS FOAM SUBSTRATE



OBS. 07 — ALGAE BLOOM

OBS. 01 — BIOFILM ON SMOOTH SURFACE
The field study started with a walk to see where biology actually grows on everyday urban materials. Looking at flat walls made of smooth metal or glass, there is almost nothing green. Rain hits these surfaces and sheets off immediately. Gravity pulls the water away too fast for anything complex to anchor or grow. Where growth does appear on these smooth vertical faces, it is limited to a biofilm.

OBS. 02 — CORRUGATED GEOMETRY
But looking closer, life finds specific shortcuts through physical structure. On a sloped metal siding, green patches appear seemingly at random. However, the geometry dictates the pattern, the biofilm clusters on or directly next to the higher peaks of the corrugation. Algae and lichen growth is also concentrated along the horizontal seams where sheets overlap.

OBS. 03 — LUSH GROWTH ON CURB
Moving down to ground level, the scene changes completely. Thick mats of lush moss cover flat concrete curbs that separate lawn from the path; here, rainwater seeps into the soil and is retained longer during dry periods. This prolonged exposure creates an environment beneficial for bryophyte growth. Lifting these ground-level moss patches up reveals one final important detail. Underneath the green layer, there is always a trapped layer of sand, dust, and dirt.

OBS. 04 — POROUS FOAM SUBSTRATE
Moss growth does not limit itself to natural or mineral surfaces. An abandoned bike seat creates the perfect conditions to host it, not on the surface, but in the porous foam on the inside of the bicycle seat. Where the plastic cover has split open, thick moss grows straight out of the soft foam underneath, because the foam holds water far better than any smooth surface could. You can see a clear separation between the sterile, impervious plastic shell and the absorbent foam inside.

OBS. 05 — GROWTH BETWEEN THE CRACKS
On this sloped wall with a very rough, textured surface, there was an abundance of life. Almost the entire surface showed visible growth, but there was a clear distinction based on orientation. The sun-facing side had a thin layer of algae, which tolerates drier, more exposed conditions and does not need much moisture. The shaded side, by contrast, held chunks of moss packed into the grooves, since that orientation means less evaporation and less harsh UV, keeping the surface damp enough for moss to thrive.

OBS. 06 — ECOLOGICAL SUCCESSION GRADIENT
On less textured surfaces, biological growth was mainly limited to the first meter from the ground. This is because spores are blown against the base of the wall and collect there, and because moisture stays higher close to the ground. Higher up the wall you can see lichen instead, which seems less limited by height and able to survive on far less moisture, holding on where the surface dries out quickly and moss cannot establish.

OBS. 07 — ALGAE BLOOM
Where there is an overabundance of water and a steady supply of moisture, algae thrive. At the end of the downspout, where the water hits the wall, a brightly coloured patch of green is concentrated around the point of impact. The growth tracks the wettest line exactly, fading out as the surface dries with distance from where the water lands.

1.3.1 DESIGN PRINCIPLES

The growth patterns reduce to three things a bioreceptive surface has to do. Each is drawn directly from the field observations and each translates into a concrete feature of the tile.

1 CATCH AND SLOW WATER

Smooth glass and metal stayed bare because rain ran off before anything could use it, while growth appeared wherever a ledge, seam or peak interrupted the flow and briefly held water. The surface should therefore break runoff rather than shed it, using raised ridges and recessed valleys that catch rain and lengthen the time it stays on the face.

2 STORE IT IN THE MATERIAL

On the torn bike seat, moss grew only from the absorbent foam and never from the sealed plastic beside it, because porous material holds water between rain events where a smooth surface cannot. The collecting parts of the tile should therefore be porous and unglazed so they absorb and buffer water, while the exposed peaks can be glazed to shed cleanly. The boundary between glazed and porous follows the boundary between the parts that shed water and the parts that retain it.

3. SHADE IT AND TRAP ITS SUBSTRATE

Shaded grooves held thick moss while sun-facing faces carried only thin algae, and every moss mat sat on a trapped layer of sand and dust. Shelter slows evaporation and captured debris becomes the bed an organism roots in. The surface should therefore use deep, self-shading grooves with rough floors and small lips that hold both moisture and fine sediment in place.



Fig. 1.7 — Bryophyte growth between the cracks of a rough surface.

01 CATCH AND SLOW WATER



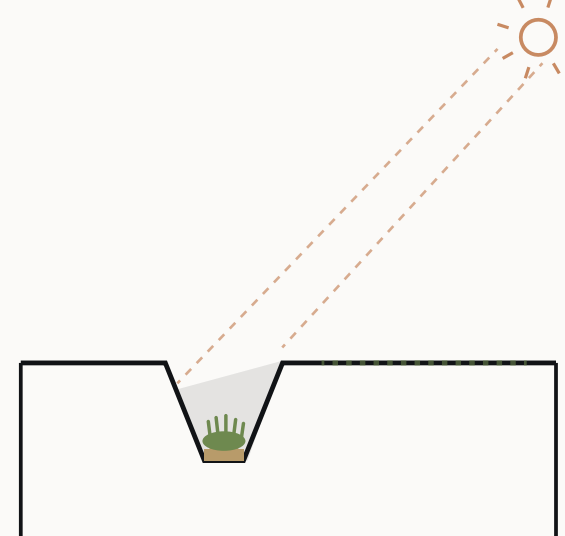


Fig. 1.8 — Design principles for bioreceptive surfaces

02 STORE IT IN THE MATERIAL



03 SHADE IT · TRAP SUBSTRATE



design principles - state of the art - living walls - water retention

1.4

State-of-the-Art Review

1.4.1 SOIL-BASED GREENING SYSTEMS AT THE GREEN VILLAGE

The most established approach to facade greening uses vascular plants grown in substrate-filled panels or cassettes mounted to the facades. These systems, known as green walls, are entirely severed from ground-level moisture and rely on an integrated water supply to keep plants alive. This dependency on irrigation is the main constraint: without the continuous water delivery, substrate volumes are too small to sustain plant life through dry periods.

Three systems currently tested at the DreamHus field lab at The Green Village, TU Delft, illustrate this constraint and the different strategies employed to address it. Together they represent the current state from passive and semi-passive soil-based facade greening in the Dutch context.

WEVERLING/ MOBILANE/RAINMATE

The most technically layered system is a collaboration between three parties. Mobilane contributes the LivePanel modular cassette architecture: individual plant units slot into horizontal aluminium gutter profiles that double as water reservoirs. Plants draw moisture continuously from this buffer through a capillary microfibre textile embedded in the cassette base, with no pump required during normal operation (Archello, n.d.). Weverling Groenproviders contributes horticultural expertise, selecting plant species and substrate compositions optimised for capillary uptake and drought tolerance (Green Village, 2025). The third component, Rainmate by Blue Innovations, is a passive rainwater distribution pipe attached to the facade that captures roof runoff and redistributes it by gravity across multiple floors, making pump-free irrigation of tall facades possible without a mains water connection (Dunea, 2023).

The combined system is being tested at The Green Village to establish how long the passive water buffer sustains plant life without any active water addition.



Fig. 1.10 — Mobilane Modular Cassette Architecture with Passive Rainwater Distribution

A site visit conducted after a five-day heatwave found the Mobilane-based installation in the best visible condition of any system on site: dense, continuous cover with high species diversity and limited desiccation. Whether this reflects resilience through the heatwave or pre-existing differences in establishment cannot be determined from observation alone, but the presence of an internal water buffer is the most plausible differentiating factor.

LEAFY

Leafy is a Dutch start-up developing lightweight, bio-based facade panels in which plants grow and is connected to a rainwater collection system. The system also relies on passive rainwater capture via panel geometry and a living substrate. Two generations are installed at DreamHus: a first-generation Accoya-wood frame design, and a newer second-generation panel developed with Bambooder and NPSP using bamboo fibre composite, optimised substrate, and an updated plant palette informed by a local flora and fauna scan (The Green Village, 2025).



Fig. 1.9 — Leafy bamboo composite panels

At the time of the site visit, the two Leafy generations showed a visible divergence in plant condition. The bamboo composite panels had desiccated vegetation spread across the cassette positions. The Accoya-wood panels showed different growth patterns, most lush vegetation was located in the bottom cassettes. This created a growth pattern forming a gradient from lush, green vegetation to dry, desiccated twigs. Whether these differences are due to gravity's effect on the water distribution or due to species-specific drought tolerance cannot be concluded from a single observation. What the pattern does suggest is that in the absence of a water buffer, survival is not guaranteed.



Fig. 1.11 — Leafy Accoya-wood panels

The contrast between Leafy and Mobilane reflects a deliberate design philosophy trade-off rather than a difference in build quality. Leafy eliminates infrastructure at the cost of transferring the entire water management burden onto the biology and substrate. Whether a fully passive approach can be made viable in a climate with increasingly frequent dry spells is precisely the question the DreamHus testing is set up to answer. 15

CRITIQUE FROM LITERATURE

The broader academic literature on living wall systems reinforces the limitations visible in the field. A widely cited review by Manso and Castro-Gomes (2015) identifies high installation costs and continuous maintenance burdens as the principal barriers to the wider adoption of living walls. A systematic review of European installations by Gunawardena and Steemers (2020) further highlights these maintenance challenges, noting that management authorities consistently struggle with plant stress and the long-term sustainability of the active infrastructure.

More broadly, an investigation in *The Architects' Journal* documented a pattern of high-profile living wall failures in the UK, such as the complete die-off of a £100,000 installation in Islington and the failure of an eight-year-old wall on an M&S store in Newcastle. These failures were directly attributed to irrigation system breakdowns and the gap between the ecological ambition of the installation and the inadequate maintenance infrastructure actually provided after handover (Ing & Jessel, n.d.). Across these sources, a consistent critique emerges: soil-based greening systems promise passive ecological performance but depend heavily on active, mechanical maintenance to deliver it.

Soil-based greening systems, regardless of their technical sophistication, share a common biological constraint: vascular plants require continuous root-zone moisture and cannot survive extended desiccation. Systems that manage this through reservoirs and capillary infrastructure perform well but carry a significant engineering overhead and fire risk when they dry out (Ing & Jessel, n.d.).

1.4.2 PASSIVE, CEMENTITIOUS BIORECEPTIVITY; RESPYRE

The initial commercial approach introduced by firms like Respyre focused on applying a homogeneous, porous concrete plaster coating directly onto existing urban facades. According to a corporate feature by Secrid (2023), these early-stage configurations relied on a temporary, top-down external pipe assembly mounted to the roofline to provide the constant moisture required for moss spores to germinate and establish a dense layer within a 10 to 12-week window. Following this initial phase, the plumbing infrastructure was intended to be removed, leaving the biological layer to sustain itself autonomously on natural rainfall.



Fig. 1.12 — Respyre Porous Concrete Coating Utilizing Temporary Irrigation Infrastructure



Fig. 1.13 — Respyre surface texture close-up



Fig. 1.14 — Respyre testing set-up at the Green Village

This early configuration was marketed on the premise, which they mention in detail on the TU Delft website, that bryophytes are "fantastic at retaining moisture" to passively lower urban microclimatic temperatures via continuous evaporation (TU Delft, 2023). However, this assumption contradicts established bryophyte physiology. This is because non-vascular plants lack protective cuticles and internal water regulation mechanisms, their cellular moisture evaporates completely passively (Jakubovskis, 2023). Meaning that they are inherently incapable of prolonged moisture retention. These findings, combined with the fact that the moss was only actively watered for a limited period, compromises the prototype's ability to achieve its cooling goals.

Additionally, testing by Veeger et al. (2026) demonstrated that when a porous cementitious substrate lacks a substantial internal water containment matrix, the moss layer experiences rapid "flash drying" during peak summer heatwaves. Once the moisture is depleted, the organisms enter metabolic dormancy and cease evaporation entirely.

Furthermore, Veeger et al. (2026) revealed that bone-dry, dormant moss layers under direct sunlight caused localized surface temperature spikes between +1.5°C and +4.1°C above bare concrete baselines. At night, this dry biological carpet acts as an insulating barrier that stifles radiative heat dissipation, trapping daytime thermal mass within the building envelope and yielding substrate temperatures +3.0°C to +3.8°C hotter than bare concrete structures. This suggests that in dormant states, these facades may inadvertently exacerbate the UHIE, exemplifying the importance of having a continuous supply of moisture.

Recent updates on their website indicates that the Respyre is adapting to these limitations. Visual documentation published on Respyre's digital platform suggest an apparent pivot away from single-layer, surface-applied coatings toward multi-component panel systems, such as their updated Mosscrete™ cavity walls and insulated assemblies (Respyre, n.d.). Promotional animations of these systems depict a continuous fluid layer, seemingly an integrated liquid water buffer, flowing directly behind the porous concrete facade panel face. While the internal mechanics of this configuration cannot be definitively verified without public technical datasheets, this design iteration strongly signals that material surface porosity alone is insufficient to sustain bryophytes during urban droughts. Exemplifying the need for an integrated water reservoir, to keep the evaporation effective.

02

AIM & RESEARCH PLAN



Project Objective

The objective is to develop a bioreceptive ceramic facade that supports living moss on its surface and keeps that moss cooling through the periods when cooling matters most. The problem is already set out: moss cools only while it is evaporating, but it cannot store its own water, so in a heatwave an unprotected layer dries out, falls dormant, and becomes a thermal liability rather than a benefit (Veeger et al., 2026).

The project's response is to move the job of water management from the plant to the material. By engineering the porosity of a 3D-printed ceramic so that it harvests rainwater, stores it inside its own structure, and feeds it back to the surface by capillary action, the tile is designed to act as an artificial vascular system that keeps the moss supplied between rain events. The main mechanism is rain capture and storage, which becomes more important as rain events grow shorter and more intense; overnight dew is treated as a helpful secondary input rather than something the design depends on (Stohl et al., 2023).

Success is practical rather than tied to a single number: the system works if, using only the water it captures and with no external watering, it stays wet for long enough to keep the moss active through the dry spells of a peak summer. Because a full summer dry period cannot be reproduced in the available time, this is judged by measuring how much water the material holds and how slowly it releases it under continuous worst-case drying, benchmarked against an unbuffered substrate. That gives a conservative, lower-bound estimate of how long the surface could sustain itself in the field.

The core of the work is the material and surface engineering, validated experimentally at the micro and meso scales. Scaling the validated tile into a full facade module, with mounting and water distribution across a wall, is developed as a conceptual application: a demonstration of how the tile could extend into a facade, and of the constraints a production design would still need to resolve.

2.2

Research Framework

From Micro to Macro

This chapter establishes the research methodology for developing the bioreceptive facade system. The project follows a Research-through-Design (RtD) approach. This method uses iterative prototyping and physical testing to generate design knowledge. The framework is structured across three interconnected phases, each representing a distinct research scale with specific inputs, experimental methods, and design outputs. Water management drives the framework. This emerged directly from the evidence. Field observations confirmed that spontaneous moss colonization in urban environments is tightly bound to zones of consistent microclimatic dampness. Expert interviews at the Green Village identified water availability as the critical variable that most bioreceptive designs fail to address. And the Respyre case study demonstrated that inadequate moisture retention leads to systemic failure during extended dry periods, an outcome that proved to be worse than no biological layer at all. (Veeger et al., 2026)


This focus is also inseparable from the primary climate goal. For the moss layer to actively lower surrounding temperatures, it must keep evaporating. The moment it dries out and enters dormancy, the cooling stops. The system must stay consistently wet through the hottest days, precisely when the urban heat island effect is most severe. Every phase in this framework is therefore tuned to one thing: controlling how water moves through the system: absorbed at the micro-scale, retained at the meso-scale, and distributed at the macro-scale.

2.2.1 RESEARCH QUESTIONS

MAIN RESEARCH QUESTION:

How can 3D-printed porous ceramics be used to design modular bioreceptive facade tiles that sustain living moss as a passive evaporative-cooling surface ?


01 MICRO-SCALE



MATERIAL

How can the ceramic body be engineered to reach the target bioreceptive porosity while remaining frost-resistant and chemically suitable for colonisation?


02 MESO-SCALE



MODULE

How does the printed surface geometry affect the way water is captured and retained at the tile surface, and which geometry shows the most promise for prolonging near-surface moisture?

03 MACRO-SCALE



SYSTEM

How can the validated tile be integrated into a facade module, and what does a prototype reveal about its cooling and its capacity to sustain moss?

Fig. 2.2 — Research questions diagram

2.2.3 THE THREE PHASES

Phase 1: Micro-Scale (Material Synthesis & Validation)

The first phase addresses substrate engineering through empirical laboratory testing. The research begins by isolating variables to test the singular effects of different additives and firing temperatures on total apparent porosity. These individual performance curves are then mapped and combined.

To ensure survival in cold climates, the final recipe is calculated using the Maage Factor to target frost resistance while maintaining the target porosity of bioreceptivity of 25% to 35%. Finally, an aqueous suspension test is conducted to verify a neutral pH value. This confirms the substrate will not leach harmful compounds, ensuring foundational biological viability.

Phase 2 Meso-Scale (Geometric Exploration & Climate Testing)

The second phase translates the optimized clay recipe into individual physical forms using Liquid Deposition Modeling (LDM). Designs are developed based on surface geometries found in literature. This phase explores how different 3D-printed toolpath textures and macro-geometries alter how water moves, pools, and spreads across the tile face.

Rather than seeking a single mathematical optimization, this phase operates as a comparative exploration. Multiple distinct design iterations are printed and placed into a climate chamber. These prototypes are tested directly against each other to evaluate their real-world sorptivity and hydration retention over time. The goal is to identify which geometric features are most effective at prolonging moisture levels.

Phase 3: Macro-Scale (Systemic Adaptation & Assembly)

Phase Three is a synthesis and validation cycle rather than a controlled study. It combines the validated recipe and the selected geometry into an integrated module, tests whether that module supports moss and delivers cooling, and develops the scaling into a facade as a conceptual application. Its results are therefore first functional outcomes and feasibility demonstrations, not optimised specifications.

03

PHASE 1 MICRO-SCALE MATERIAL SYNTHESIS

Unplanned bioreceptivity — field research, observations, and the state-of-the-art in bioreceptive facade design.

Fig. 3.1 — Microscopic Imaging of Diatomaceous Earth (400x)

3.1

Introduction

The macro geometry, defined by the general shape of a facade tile, is vital for creating biological niches and localized microclimates that shield organisms from the elements (Mustafa et al., 2021). However, the success of these macro-features ultimately depends on the micro-scale properties of the material itself (Rotondi et al., 2024). One of the most important factors for a bioreceptive surface is the ability to absorb and store water: an essential building block of life (Parveen & Tariq, 2023). This is achieved by increasing the material's porosity. Pores function as capillary channels, drawing water against gravity through the structure, while also serving as internal basins to store moisture during dry periods (Rotondi et al., 2024; Veeger et al., 2021)

Furthermore, it is important to create a rough surface with cavities and ridges where cells can accumulate and rhizoids, the "root-like" appendages of mosses, can securely attach (Veeger et al., 2026). To create these conditions, we must engineer the ceramic on a micro-scale. While ceramics can be inherently porous if fired under specific conditions, a technical trade-off exists: the lower the firing temperature, the more porous the material remains, but it also experiences reduced flexural strength and overall structural integrity (Rotondi et al., 2024).

Another critical issue with increased porosity is environmental durability. Pores between 0.1 μm and 3.0 μm are considered a "critical" size range for frost damage, as they allow water to enter but do not always accommodate its expansion (Stryszewska & Kanka, 2022). This risk can be mitigated by creating a bimodal, hierarchical pore structure (Li et al., 2013). When a material possesses larger pore sizes outside of this critical range (greater than 3.0 μm), they function as expansion chambers for ice, preventing internal crystallization pressure from damaging the structure (Stryszewska & Kanka, 2022).

In addition to moisture retention, bioreceptivity is heavily influenced by the chemical environment of the material, particularly its surface pH (Parveen & Tariq, 2023). While traditional construction materials like standard Portland cement exhibit high alkalinity (pH > 12) that inhibits most biological colonization (Parveen & Tariq, 2023), ceramics and clay-based materials are generally less alkaline and inherently more compatible with a wider range of organisms (Parveen & Tariq, 2023). Raw clay mixtures typically begin with a neutral to slightly alkaline pH (around 7.0–8.0) (Rotondi et al., 2024). During the firing process, depending on the peak temperature and specific mineral composition, the resulting ceramic can either maintain a near-neutral state (e.g., pH 6.8–7.0 for certain low-fired bodies) or experience an increase in alkalinity up to a maximum pH of 10 (Rotondi et al., 2024; Tietz, 2021).

The primary objective of Phase One is to create a functioning ceramic recipe that achieves this interconnected hierarchical structure. By utilizing spent coffee grounds (SCG) to create macro-pores (Maciel et al., 2023; Rotondi et al., 2024) and diatomite (diatomaceous earth) for its dual-pore structure consisting of internal micro-pores within frustules and macro-pores between particles (Jiang et al., 2019). The goal is to reach a target porosity of 25% to 35% (Mustafa et al., 2021). This phase identifies the optimal balance between bioreceptivity and environmental resilience.

3.2.4 SACRIFICIAL ADDITIVES

A common technique to increase the porosity of a ceramic is the inclusion of a sacrificial additive, which often consists of organic materials derived from agricultural or industrial waste streams (Salleh et al., 2021). These organic materials are mixed into the wet clay body and act as pore-forming agents. During the firing process, as the kiln reaches temperatures around 600 °C, the organic matter undergoes pyrolysis and completely burns away, leaving behind empty cavities or "ghost porosities" in the ceramic (Parveen & Tariq, 2023). As the temperature continues to increase towards the peak firing temperature, the surrounding clay matrix begins to sinter and partially vitrify, fusing the ceramic particles together to create a hard, consolidated skeleton around the newly formed pores (Liu, 2011).

Not every additive will produce the same pore shape and size, which directly affects the material's mechanical strength and its ability to absorb and retain water (Hammel et al., 2014). For example, materials with a high particle form factor can produce round pores, whereas others may create oriented or elongated pores (Nigay et al., 2017). The choice of additive also significantly impacts the processability of the clay mixture during 3D printing (Rotondi et al., 2024). Granular or spheroidal additives facilitate smooth extrusion, whereas larger or highly fibrous materials are prone to clogging the printer nozzle (Rotondi et al., 2024). Furthermore, highly absorbent organic materials can retain excess water within the paste, altering the drying rate and potentially leading to uneven shrinkage or cracking during the fabrication and sintering processes.

CARDBOARD PULP

A common sacrificial additive used to increase porosity is cardboard pulp, which Hadi and Al-Hussien (2017) note is economical, non-toxic, and easy to burn. Furthermore, Rotondi et al. (2024) explains that cardboard mainly consists of tubular cellulose fibers, resulting in pores that mirror this shape. In their study on bioreceptive surfaces, the authors state that this balances water permeability and strength. However, the researchers eventually observed that the continuous channels cause the material to dry out too quickly. Consequently, Rotondi et al. recommend moving away from fiber-based additives, which clog 3D printer nozzles, in favor of spheroidal coffee grounds.

SAWDUST

Sawdust is also often added as an effective organic sacrificial additive for enhancing the bioreceptivity of 3D-printed clay composites. (Sochůrková et al., 2024) To prepare the material for smooth extrusion in additive manufacturing, it is processed through a 0.5 mm sieve before being mixed with water and wet clay

However, there is a critical physical constraint when working with this specific additive: a maximum of 10% sawdust can be added to the clay mixture (Sochůrková et al., 2024). If the concentration exceeds this threshold, the clay rapidly loses its necessary plasticity and takes on a weak, "sandy" character. This loss of plasticity impairs both the smooth extrusion of the paste and the structural stability of the printed geometry.

SPENT COFFEE GROUNDS

Spent coffee grounds (SCG) are a fairly new ingredient in the circular economy. Due to the granular morphology of SCG particles, it is possible to obtain a homogeneous distribution within the clay, which is important for achieving consistent material behavior in 3D printing. This fine-grain size also prevents the structural clogging often associated with cellulose-based fibers (Rotondi et al., 2024). Beyond fabrication benefits, the granular nature of the waste contributes to a high degree of surface roughness and the formation of micro-grooves. These features function as shaded micro-niches that improve the initial adhesion of cells and the secure attachment of rhizoids (Rotondi et al., 2024). Hartanto and Janasthi (2023) note that a percentage range of 5% to 15% provides optimal porosity without catastrophic loss of structural integrity, making it ideal for creating functional, self-watering ceramic planters.

STARCH

Unlike other organic additives, starches can swell and gelatinize in water, causing it to act as both a pore-forming agent and a binding agent during the consolidation process (Abdul Rahman et al., 2021; Li et al., 2013). When starch burns out, the resulting pores become interconnected, which practically dictates the ceramic's water permeability and accessibility to fluids (Research, Society and Development, 2021). For example, a study demonstrated that adding 20% starch by weight to a kaolin clay matrix increased the total porosity of the support from approximately 27% to 46%, and successfully enlarged the average pore size from 0.5 μm to 1.4 μm (Bouzerara, Harabi, & Boudaira, 2009).

Corn starch yields finer pore networks and higher mechanical strength, whereas potato starch generates larger pores and greater total porosity (Pabst et al., 2017). Furthermore, modified cassava starch has been shown to optimize the material's overall mechanical response and fracture strength, particularly at high temperatures (Tomba Martinez, Sandoval, & Camerucci, 2012).

3.2.5 DIATOMITE

Diatomite, or diatomaceous earth, is an inorganic material formed from the fossilized silica shells of single-celled diatoms. Characterized by high amorphous silica content and low thermal conductivity, it is frequently sourced as a saturated waste byproduct from the brewing industry for use in sustainable ceramic production (Silva et al., 2022).

Silva et al. (2022) note that this waste serves as a highly effective inorganic pore-forming agent in the production process of highly porous ceramic foams. Adding diatomite reduces the material's bulk density and enhances its insulating properties. Furthermore, the material is remarkably efficient; while organic additives often require high volumes, adding as little as 1% to 5% diatomite by weight can produce bricks with up to 39.4% porosity.

If diatomite is simply mixed into a wet clay body, the increase in porosity relies entirely on the intrinsic, natural structure of the fossilized diatom shells; this generates microscopic pores typically measuring between 2 and 5 μm (Processing Methods for the Preparation of Porous Ceramics, 2014).



Fig. 3.4 — Micrograph of diatoms (diatomaceous earth). Retrieved from <https://pbs.twimg.com/media/EVnWphU4AAxLHd.jpg>

3.2.6 FROST DURABILITY

Although high porosity is essential for bioreceptivity, materials within the target range of 25–35% total porosity are more susceptible to internal ice crystallization pressure. Frost durability can be predicted using the Maage Factor (Fc):

$$F_c = (3.2 / PV) + 2.4 * P_3$$

Where PV is the total pore volume and P3 is the percentage of pores larger than 3.0 μm. A material is considered frost-resistant when $F_c > 70$ (Stryzewska & Kanka, 2022). This formula reveals that frost resistance is not solely a function of total porosity, but critically depends on the internal pore size distribution. Larger macropores (> 3.0 μm) act as expansion chambers where freezing water can safely expand, relieving crystallization pressure on the surrounding ceramic matrix. This creates a clear design constraint: achieving the bioreceptive porosity target must be accompanied by a sufficient proportion of macropores to satisfy the Maage threshold, making a hierarchical, bimodal pore structure a structural necessity rather than an optional feature.

Due to the unavailability of Mercury Intrusion Porosimetry (MIP) to empirically measure the exact pore size distribution within the fired samples, a different approach was used to calculate the distribution of pore sizes based strictly on the raw particle size and burnout behavior:

- **Macropore Origin (> 3.0 μm):** It is assumed that all porosity generated by the combustion of Spent Coffee Grounds (SCG) results in large, stress-relieving macropores. This is justified by the large, granular physical size of raw SCG particles.
- **Micropore Origin (< 3.0 μm):** It is assumed that the inherent porosity of the sintered clay matrix, as well as the voids created by the fine-grained secondary additives (cornstarch and diatomite), exclusively make up the microscopic capillary network.

It should be noted that in practice, a proportion of the pores generated by corn starch and diatomite will inevitably exceed the critical frost damage threshold of 3.0 μm, functioning as additional expansion chambers. This means the binary classification applied here likely underestimates the true frost resistance of the material, introducing a conservative margin into the Maage Factor calculation. While this simplification does not fully capture the complex pore size distribution of a fired ceramic body, it provides a reliable lower-bound estimate of frost resistance without requiring mercury intrusion porosimetry (MIP).

3.2

Literature review

3.2.1 CLAY TYPOLOGY

Engineering the porosity of a ceramic body begins with selecting the appropriate clay. Clays are naturally occurring materials characterized by a diverse range of mineralogical and chemical compositions (Rotondi et al., 2024). All these minerals interact during the firing process, directly influencing both the structural integrity and the resulting porosity of the final piece (Cultrone et al., 2004; Jordan et al., 2008). While variability exists within clay types based on their mining origin and additives, they are generally categorized into three main types; earthen-, stone- and porcelain

Earthenware is the most naturally porous of the three types. It is typically bisque-fired at low temperatures and is widely used for architectural applications, garden pots, and roof tiles due to its excellent ability to retain and release water through capillary action (Rotondi et al., 2024). While its high water absorption (typically 5–15% or higher) is beneficial for bioreceptivity, the lack of full vitrification makes it softer and more prone to chipping or spalling from internal ice expansion in exterior freeze-thaw conditions (Stryszewska & Kanka, 2022).

Stoneware clays are usually white, grey, or yellowish and are significantly stronger than terracotta. When fired at their intended temperature (typically between 1000 °C and 1300 °C), they become dense and stone-like (Flynn & Stachurski, 2006). At these elevated temperatures, the clay matrix undergoes partial vitrification and significant densification, resulting in very low open porosity and water absorption rates, which typically drop from around 19% down to between 0.5% and 3% (Flynn & Stachurski, 2006).

Porcelain is fired at the highest temperatures, typically between 1200 °C and 1400 °C (Ochen et al., 2019). It is extremely hard and effectively possesses no open porosity, being officially classified as "impervious" to moisture with a water absorption rate of less than 0.5%, and often below 0.1% in commercial tile products (Ochen et al., 2019; Zanelli et al., 2004). Because this glassy phase seals off internal pores, this lack of moisture storage makes it generally unsuitable for bioreceptive applications (Rotondi et al., 2024).

When developing and engineering bioreceptive ceramics, researchers predominantly use earthenware, particularly white earthenware (Rotondi et al., 2024). Studies focusing on 3D-printed or highly porous bioreceptive facades specifically select white earthenware because this fine-grained, highly plastic material safely maintains a high open porosity (e.g., around 8% natural water absorption) when bisque-fired at low temperatures (around 960 °C to 1000 °C) (Rotondi et al., 2024). Unlike stoneware or porcelain, which typically vitrify into a liquid glass phase that seals off capillary channels, low-fired earthenware ensures the ceramic network remains highly permeable to fluids (Rotondi et al., 2024).

Other researchers have successfully adapted stoneware clay for bioreceptive and environmentally responsive 3D-printed architectures by intentionally under-firing the material to prevent full vitrification. For example, in the development of "probiotic" architectural scaffolds designed to host lichen and fungal networks, researchers specifically selected 3D-printed white stoneware clay that was bisque-fired at a lower temperature of 1063 °C (Cone 04) (Beckett, 2021).




Fig. 3.2 — Open porosity by clay type.

3.2.2 POROSITY

The target porosity for a bioreceptive building material generally falls within the range of 20% to 40% to achieve the optimal balance of moisture retention required to sustain biological growth (Veeger et al., 2021; Lim & Lharchi, 2023).

A high volume of interconnected pores significantly increases the material's ability to absorb and hold capillary water. Because organisms like mosses and lichens depend entirely on their surrounding environment for hydration, this internal water reservoir keeps them active during dry periods (Guillitte, 1995; Proctor et al., 2007). Pores also function as physical shelters at the micro-scale, trapping windblown spores, atmospheric dust, and nutrients while protecting pioneer microorganisms from UV radiation, wind detachment, and rapid desiccation (Cruz & Beckett, 2016; Gorbushina, 2007).

There is a fundamental trade-off in bioreceptive design: while higher porosity improves biological colonization and water absorption, it simultaneously reduces mechanical integrity. Every pore acts as a stress concentration point where cracks can initiate under load, so compressive strength, elastic modulus, and fracture toughness all decrease as porosity increases (Veeger et al., 2021).

3.2.3 PORE SIZES & BIORECEPTIVITY



01
CAPILLARY & WATER RETENTION

0.1–3.0 μm

Pores in this range absorb and hold water against gravity through capillary suction, a foundational requirement for bioreceptivity (Guillitte, 1995). This same range is also the critical danger zone for freeze-thaw damage (Veeger, Ottel   & Prieto Hoces, 2021).



02
SINGLE-CELLED ORGANISMS

10–20 μm

Single-celled microorganisms such as subaerial cyanobacteria need a pore size or surface roughness slightly larger than their own cell for optimal physical attachment and anchoring (Lim & Lharchi, 2023; Pentecost & Whitton, 2012).



03
FILAMENTOUS & LARGE ALGAE

500 μm

A surface niche scale of exactly 500 μm has been shown to optimize the physical entrapment and growth rates of these specific organisms (Lim & Lharchi, 2023).



04
MACRO-ALGAE & SMALL MOSSES

3.0 mm

Larger niche scales and groove sizes of around 3000 μm are optimal for the larger structure of macro-algae and the initial growth stages of small mosses (Mustafa, Prieto & Ottel  , 2021).



05
MOSS & LICHEN ANCHORING

> 5.0 mm

Resilient cryptogamic covers require macro-features larger than 5 mm to provide physical anchoring, shield organisms from wind detachment, and offer micro-shading (Cruz & Beckett, 2016; Mustafa, Prieto & Ottel  , 2021).



06
FREEZE-THAW COMPENSATION

> 3.0 μm

Pores larger than 3.0 μm are safe from frost: they act as compensation chambers where freezing water can expand without cracking the material (Veeger, Ottel   & Prieto Hoces, 2021).

BIORECEPTIVE PORE / FEATURE SIZE · LOGARITHMIC SCALE



Fig. 3.3 — Bioreceptive pore and feature sizes.

3.2.7 PH

Unlike conventional concrete, which exhibits a highly alkaline and hostile surface pH (12.0 to 14.0), ceramics and clay-based materials are inherently much more compatible with biological colonization (Parveen & Tariq, 2023). Raw clay mixtures typically possess a neutral to slightly alkaline pH ranging from 7.0 to 8.0, which naturally aligns with the optimal conditions required for pioneer soil microbes, algae, and lichens to establish (Crawford et al., 2023); (Rotondi et al., 2024). For more advanced cryptogamic covers like mosses, the targeted pH ideally drops slightly lower, as mosses strongly prefer a slightly acidic to neutral environment ranging from pH 5.0 to 6.5 (Innovative Construction Technologies, n.d.).

However, the final chemical bioreceptivity of a ceramic material is heavily influenced by its manufacturing process. During firing, especially at high temperatures, the clay undergoes mineralogical transformations that can increase the product's surface pH up to 10 (Rotondi et al., 2024). Additionally, the inclusion of alkaline fluxing agents or high amounts of calcium oxide (CaO) in the ceramic mixture can react with water to generate hydroxide ions, further elevating the surface alkalinity (Kinoshita et al., 2014).

While ceramics do not rely on the lengthy natural carbonation process that concrete requires to become habitable, their surfaces still undergo beneficial weathering over time. Once placed outdoors, early pioneer microorganisms (bacteria and fungi) settle on the ceramic and secrete organic acids. This metabolic activity physically and chemically disrupts the surface, naturally etching the material, neutralizing the pH, and releasing bioavailable trace minerals (like calcium and magnesium) that prepare the substrate for the eventual anchoring of moss rhizoids (Jakubovskis, 2024); (Parveen & Tariq, 2023). Therefore, to optimize primary bioreceptivity for moss in ceramics, designers should target lower firing temperatures and careful additive selection to maintain an initial pH as close to neutral as possible (Rotondi et al., 2024); (Crawford et al., 2023).

3.2.8 APPLIED MATERIAL CHOICES

Building on the literature above, the ceramic recipe developed in this research uses a white stoneware clay (G&S 254), spent coffee grounds, and corn starch. Diatomite was also evaluated as an alternative secondary additive.

G&S 254 is a white stoneware, but it is used well below its intended sintering range. Firing at 1100 °C, the low end of its 1000–1280 °C window, deliberately under-fires the body so it does not fully vitrify and instead keeps the open, interconnected porosity usually associated with earthenware. This is a direct response to the concern raised in the literature review, that stonewares seal their capillary network on firing: that sealing happens only at full vitrification, and under-firing avoids it, as Beckett (2021) showed with white stoneware bisque-fired at 1063 °C.

Spent coffee grounds are selected to generate the macropore fraction of the pore structure. Their granular morphology is expected to burn out completely during firing, leaving large voids above the 3.0 μm frost-damage threshold that act as expansion chambers for ice. Beyond porosity, their granular shape is compatible with LDM extrusion and contributes to surface roughness at the micro-scale, supporting cell adhesion and rhizoid attachment.

Corn starch is selected to generate the micropore fraction. Its finer burnout is expected to create pores in the capillary range, forming the internal moisture transport network that draws water upward against gravity and sustains biological activity during dry periods. Together with the SCG-generated macropores, the goal is a bimodal pore structure that simultaneously meets the bioreceptive porosity target of 25–35% and satisfies the Maage Factor threshold of $F_c \geq 70$.

Diatomite is evaluated as an alternative to corn starch as a micro-porosity agent, given its published potential for creating sub-micron pores within its fossilized frustule structure and its availability as a waste byproduct from the brewing industry.




Fig. 3.8 — Bimodal pore structure from combined additives.

sacrificial additives - bimodal porosity - research questions

3.3

Research Questions

POROSITY CHARACTERIZATION OF INDIVIDUAL ADDITIVES

RQ1 (Macroporosity):

How do varying concentrations (0–15%) of spent coffee grounds (SCG) influence the development of open macroporosity and total volume across a firing range of 900°C to 1100°C?

RQ2 (Microporosity):

To what extent do corn starch and diatomaceous earth, as independent variables, influence the formation of micropores and total open porosity within the ceramics at different firing temperatures?

OPTIMIZATION FOR BIORECEPTIVITY AND FROST DURABILITY

RQ3

What are the optimal ratios of spent coffee grounds and corn starch required to produce a dual-porosity bioreceptive ceramic that satisfies the theoretical Maage Factor of 70 for frost resistance, and at what temperature is this balance best achieved?

RQ4

What are the optimal ratios of spent coffee grounds and diatomaceous earth required to produce a dual-porosity bioreceptive ceramic that satisfies the theoretical Maage Factor of 70 for frost resistance, and at what temperature is this balance best achieved?

FINAL RECIPE

RQ5

How do varying concentrations (25–40%) of chamotte influence the porosity and total volume fired at 1100°C?

RQ6

Does the fired ceramic composite leach alkaline or acidic compounds into its moisture environment, and is the resulting solution pH compatible with pioneer organism colonization?

3.4

Methodology

3.4.1 VARIABLE SELECTION

The base clay selected for this research is G&S 254, a white stoneware with 25% chamotte and a firing range of 1000°C to 1280°C. It was chosen for its structural stability and wide sintering window. No published data exists on how this specific clay body interacts with organic or siliceous additives, making empirical characterisation a necessary first step before using it in a bioreceptive recipe. Firing temperatures of 900°C, 1000°C, and 1100°C were selected to capture the transition from a highly porous under-fired state to the onset of vitrification. Testing at 1100°C is particularly important, as it represents the lower boundary of sufficient structural sintering while potentially remaining below the point of full densification that would seal the capillary network.

SCG concentrations of 0%, 5%, 10%, and 15% were chosen to characterise macropore formation across a range that previous literature identifies as compatible with LDM printability and structural integrity. Corn starch was tested at the same concentrations to characterise its independent contribution to the finer capillary micropore network. Diatomite at the same range was evaluated as an alternative micro-porosity agent, given its published potential for creating sub-micron pores within its fossilised frustule structure.

Two tests were conducted to characterise the fired samples. Open porosity was measured using the ASTM C20-00 standard, which applies the Archimedes principle to quantify a material's capacity for water retention relative to its density. pH was measured following ISO 787-9, an aqueous suspension test that evaluates the material's chemical leaching behaviour and determines whether its moisture environment is compatible with pioneer organism colonisation. Together the two tests cover the two non-negotiable requirements for bioreceptivity established in the literature: sufficient water retention and a biologically safe pH.

3.4.2 EXPERIMENT DESIGN

To characterize the influence of individual variables on the ceramic matrix, a full-factorial design matrix was implemented. This approach ensures that every additive concentration is evaluated across the defined sintering range (900–1100 °C), allowing for an accurate porosity mapping of the material's performance and structural stability. To determine the effect on the pH of demineralised water, 3 individual ceramic samples were measured 3 times at 3 time intervals.

Porosity Mapping Experiment Design

Additive Type	Concentrations (% wt)	Firing Temperatures (°C)	Total Samples
Control	0%	900, 1000, 1100	17
Spent Coffee Grounds	5%, 10%, 15%	900, 1000, 1100	20
Corn Starch	5%, 10%, 15%	900, 1000, 1100	46
Diatomite (Normal)	5%, 10%, 15%	1000, 1100	18
Diatomite (Short Soak)	5%, 10%, 15%	900, 1000	8
Total			109

Table 3.1 — Porosity mapping experiment design.

pH Testing Experiment Design

Time Interval	Sample Material	Number of Samples	Tests per Sample	Total Tests
—	Distilled Water	1	3	3
5 Minutes	Ceramic	3	3	9
1 Hour	Ceramic	3	3	9
24 Hours	Ceramic	3	3	9
Total				30

Table 3.2 — pH testing experiment design.

PREPARATION WORKFLOW

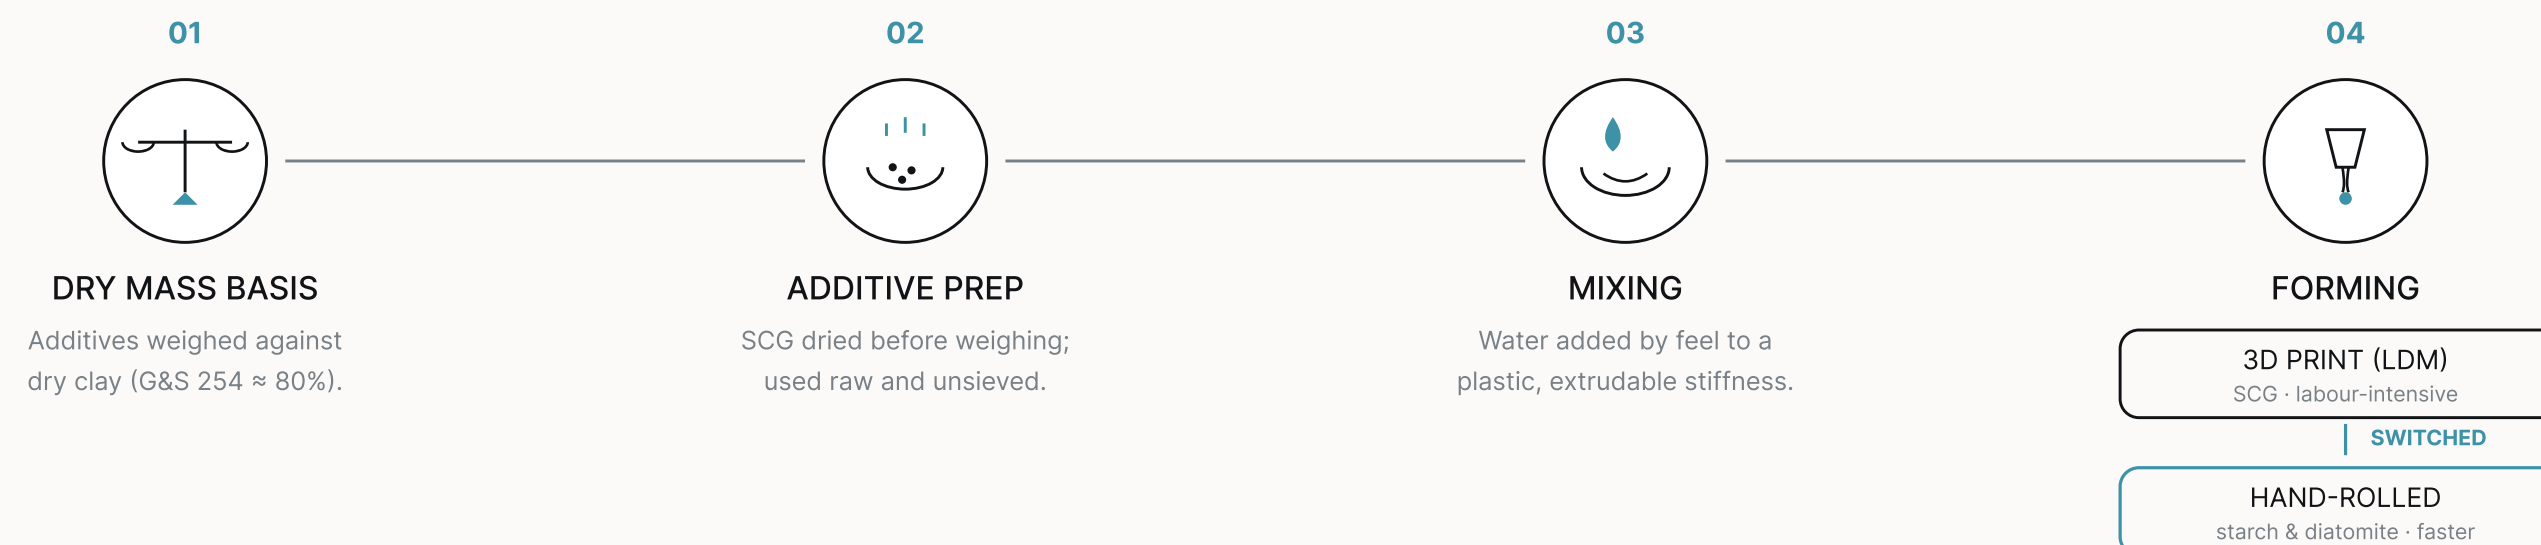


Fig. 3.9 — Preparation workflow.

3.4.3 FABRICATION AND SAMPLE PREPARATION

- **Dry Mass Calculation**: All additives were calculated based on the dry weight of the clay. The G&S 254 stoneware was assumed to have a 20% moisture content (80% dry matter); for example, a 10% additive batch utilized a mass of additive equal to 10% of the calculated dry clay mass.
- **Additive Preparation**: Spent Coffee Grounds (SCG) were dried completely prior to weighing to ensure moisture from the waste stream did not skew the additive ratios. Due to equipment availability, additives were used in their raw, unsieved state.
- **Mixing and Consistency**: Water was added incrementally to the clay-additive mixture until a plastic, extrudable consistency was achieved. While precise water-to-clay ratios were not recorded to allow for manual adjustments in workability, a consistent "stiffness" was maintained across samples.
- **Forming Methods**: Initially, coffee ground samples were fabricated using 3D-printing to test their behavior in a Liquid Deposition Modeling (LDM) workflow. However, this proved to be a labor-intensive process, leading to a change in approach. Consequently, diatomite and starch samples were hand-rolled and shaped to allow for faster iteration and to save time during testing.

3.4.4 FIRING PROTOCOLS

Samples were fired in an electric kiln under an oxidizing atmosphere. See figure

REFINEMENT FOR DIATOMITE SAMPLES

Initial results for Diatomite samples fired at 1000°C and 1100°C showed that porosity levels remained nearly identical to the baseline clay, regardless of the diatomite concentration. This indicated that the diatomite was having no measurable effect on the resulting porosity. This was likely due to the formation of a liquid glass phase during the standard one-hour soak, which vitrified and sealed the microscopic skeletal pores of the diatomaceous earth.

To potentially resolve this and preserve the material's internal structure, the firing schedule was adjusted to a minimal soaking time of 1 minute (reduced from 60 minutes) at peak temperature. This adjusted process was then applied to a new batch of samples at 1000°C and 900°C to determine if a shorter heat exposure would prevent pore-sealing while still achieving sufficient structural sintering

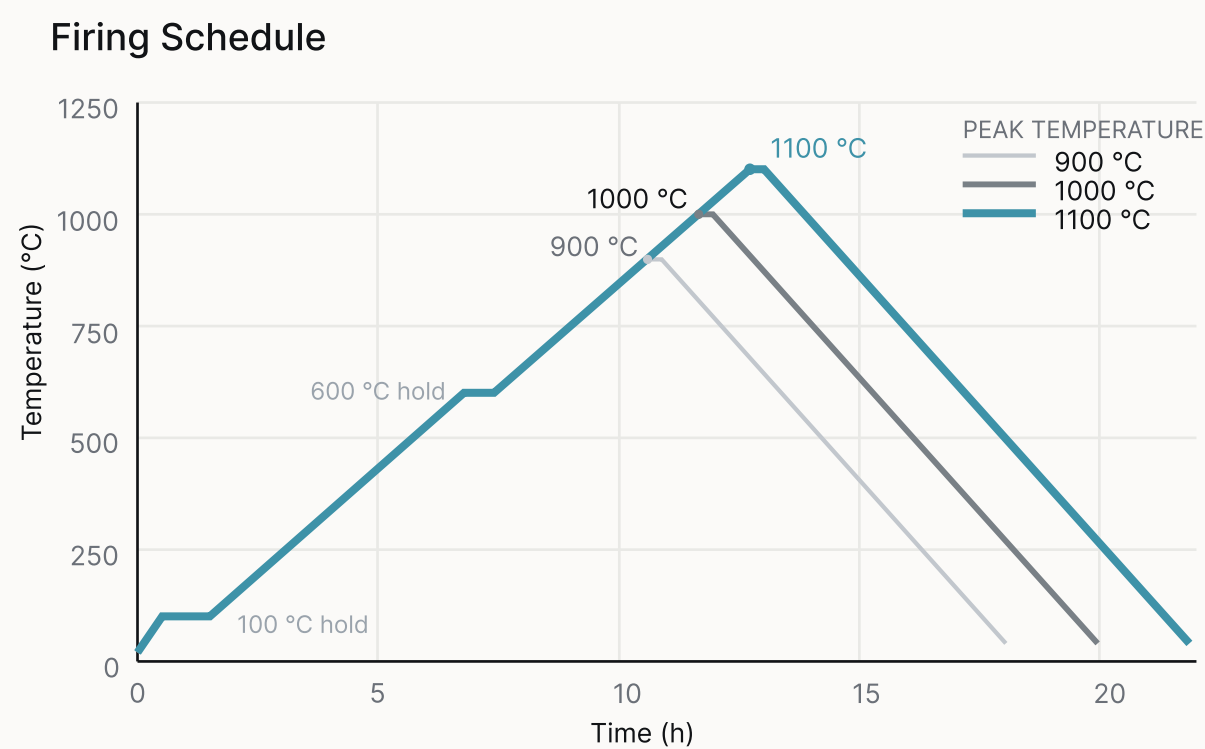


Fig. 3.10 — Firing schedule.

3.4.5 TESTING AND ANALYSIS PROTOCOLS

POROSITY EXPERIMENTAL PROCEDURE

The apparent porosity of the specimens was determined through the following steps:

1. Dry Weight (D): Specimens were oven-dried to a constant mass to establish the baseline dry weight.
2. Saturation: The samples were submerged in distilled water and boiled for two hours to displace air and ensure complete saturation of the internal capillary network.
3. Cooling: Following the boil, the specimens remained submerged as the water cooled to room temperature for a minimum of 12 hours.
4. Suspended Weight (S): While still immersed in the water, the suspended weight of each saturated specimen was recorded.
5. Saturated Weight (W): The specimens were then removed from the water. After quickly wiping away surface moisture with a damp cloth, the saturated mass in air was recorded.

Open porosity (Apparent Porosity, PV) was calculated using the following equation:

$$PV = \frac{W - D}{W - S} \times 100\%$$

PH TESTING PROTOCOL

The ceramic material was mechanically ground into a fine powder to ensure the internal core was thoroughly mixed with the exterior layer before testing. A standard 10:1 liquid-to-solid ratio was followed, creating a suspension in distilled water. pH was measured at three time intervals: 5 minutes, 1 hour, and 24 hours. Each of these three individual ceramic samples was tested three times per interval. A distilled water control was measured alongside to establish the baseline.

firing schedule - apparent porosity - spent coffee grounds

3.5

Results and Discussion

3.5.1 Apparent Porosity Characterisation of Individual Additives

EFFECT SCG CONCENTRATION ON MACRO-POROSITY

The results presented in Table 3.3 and Figure 3.11 demonstrate a clear positive relationship between SCG concentration and apparent porosity across all tested firing temperatures. As SCG content increases from 0% to 15%, porosity increases substantially at every temperature, confirming that SCG functions as an effective sacrificial pore-forming agent within the tested range.

Firing temperature has a suppressive effect on porosity, with 1100°C producing the lowest values across all SCG concentrations. This effect is already visible at 0% SCG, where the baseline clay porosity drops from 27.44% at 900°C to 16.60% at 1100°C, indicating that increased sintering at higher temperatures densifies the clay matrix independently of the additive. This suppressive effect becomes increasingly pronounced at higher SCG concentrations, as evidenced by the growing divergence between temperature curves observed in Figure 3.11.

At 900°C and 1000°C, porosity values rise steeply, reaching 64.85% and 56.10% respectively at 15% SCG. At 1100°C the same concentration yields 49.33%, reflecting the counteracting influence of enhanced matrix densification at elevated temperatures.

These concentration-temperature relationships characterize the independent contribution of SCG to macroporosity and provide the empirical basis for selecting appropriate SCG dosages in the hybrid mix designs evaluated in section 3.5.3.

Porosity by SCG Concentration

SCG (% wt)	900 °C	1000 °C	1100 °C
0%	27,44%	25,12%	16,60%
5%	34,16%	34,40%	22,83%
10%	47,65%	43,29%	35,15%
15%	64,85%	56,10%	49,33%

Table 3.3 — Apparent porosity by SCG concentration.

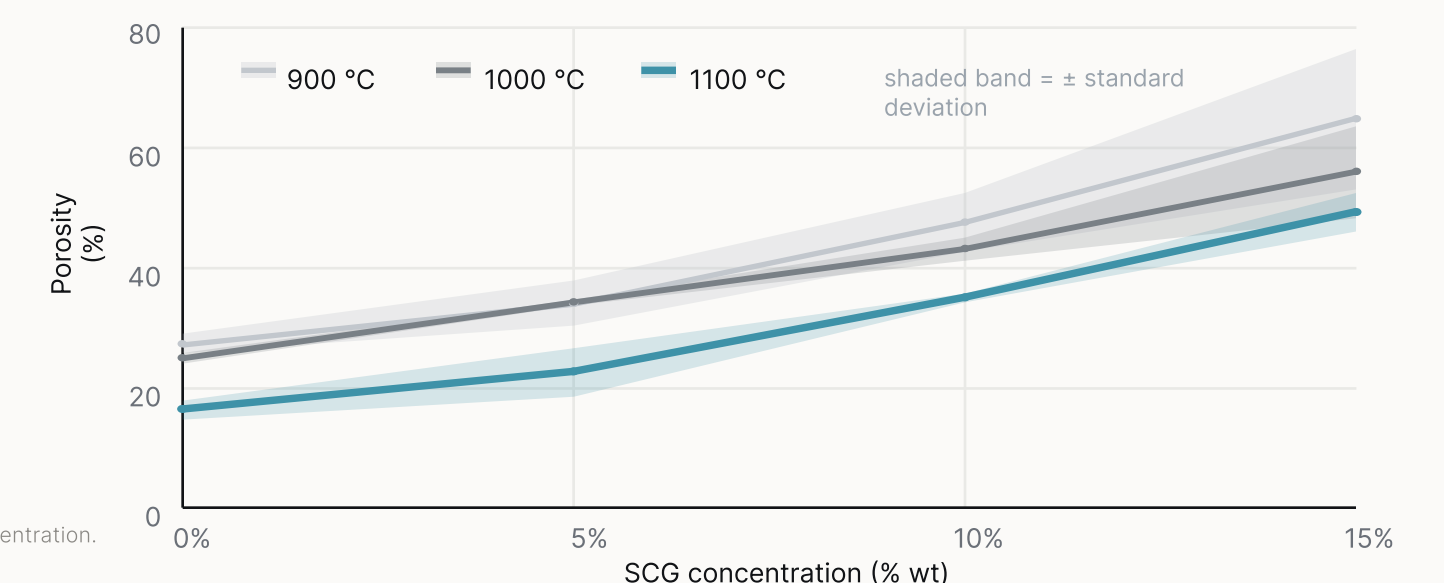


Fig. 3.11 — Apparent porosity vs SCG concentration.

EFFECT OF CORNSTARCH CONCENTRATION ON MICRO-POROSITY

Table 3.4 and Figure 3.12 present the effect of corn starch concentration on apparent porosity across the three firing temperatures. Compared to SCG, corn starch produces a notably more modest and gradual increase in porosity, rising from baseline values to a maximum of 33.31% at 900°C and 31.02% at 1100°C at 15% concentration. This restrained pore-forming behaviour is consistent with the finer pore network characteristic of corn starch burnout described in section 3.2.4.

A distinctive feature of the corn starch results is the convergence of temperature curves at higher concentrations. While the baseline clay porosity at 0% starch shows the same temperature-dependent spread observed in section 3.5.1, the three curves progressively converge as starch concentration increases, nearly meeting at 15%. This suggests that at sufficient starch concentrations, the micro-pore network generated by starch burnout partially compensates for the densification effect of higher firing temperatures.

Notably, 1100°C consistently produces the lowest absolute porosity values, yet the difference between temperatures narrows substantially compared to the SCG results. This temperature resilience makes corn starch a more stable micro-porosity agent across the tested firing range. These results characterize the independent contribution of corn starch to micro-porosity and inform the dosage selection for the hybrid mix designs evaluated in section 3.5.3.

Porosity by Starch Concentration

Starch (% wt)	900 °C	1000 °C	1100 °C
0%	27,44%	25,12%	16,60%
5%	28,99%	29,67%	23,74%
10%	29,99%	30,27%	28,85%
15%	33,31%	34,92%	31,02%

Table 3.4 — Apparent porosity by corn-starch concentration.

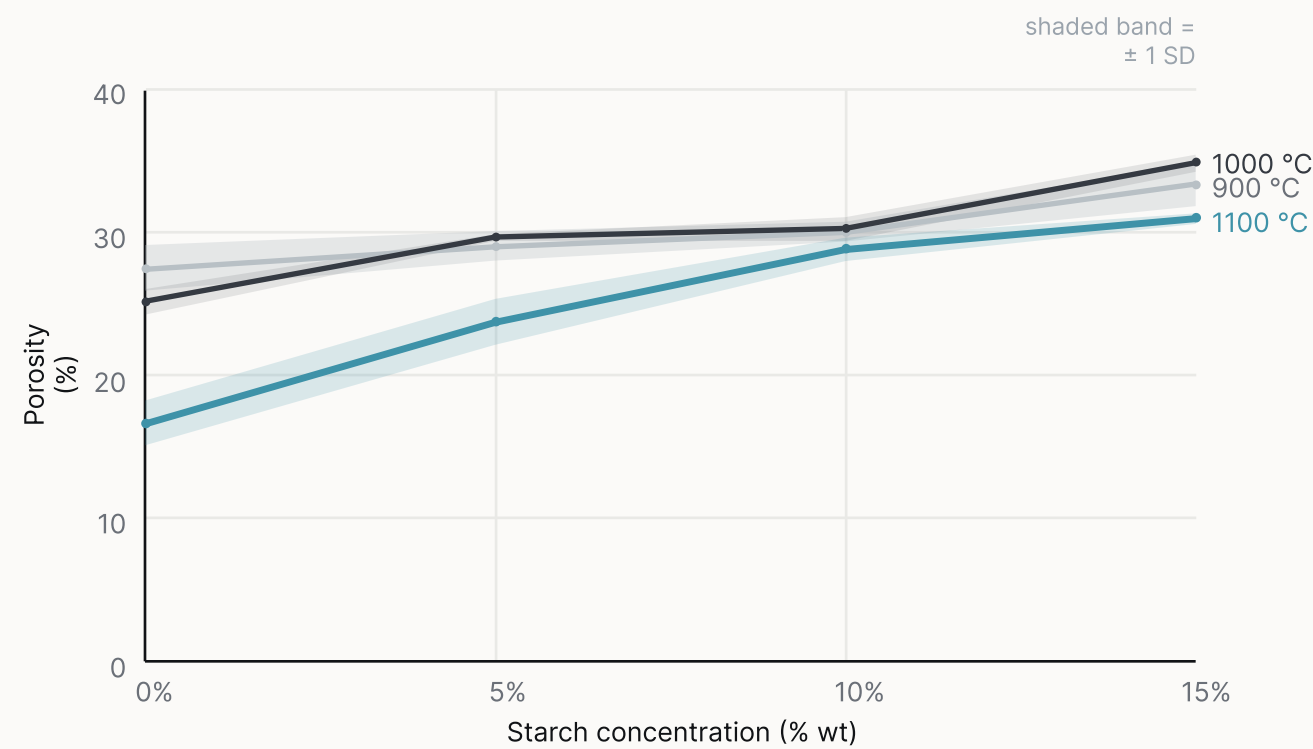


Fig. 3.12 — Apparent porosity vs corn-starch concentration.

EFFECT OF DIATOMITE CONCENTRATION ON MICRO-POROSITY

Table 3.5 and Figure 3.13 reveal a markedly different pattern from the two organic additives tested in sections 3.5.1/3.5.2. Rather than a consistent positive relationship between concentration and porosity, diatomite produces strongly temperature-dependent behaviour that diverges across the three firing conditions.

At 900°C, diatomite functions as a modest pore-forming agent, increasing apparent porosity from the 27.44% baseline to 32.71% at 15% concentration. At 1000°C, the effect is negligible; porosity remains essentially unchanged across all concentrations, ranging narrowly between 25.12% and 25.26%, indicating that diatomite contributes almost no additional pore volume at this temperature.

Most significantly, at 1100°C diatomite produces an inverse effect: porosity decreases progressively with increasing concentration, falling from 16.60% at 0% to 11.59% at 15%. This reversal indicates that at elevated temperatures the high amorphous silica content of diatomite undergoes vitrification, generating a glassy phase that actively seals existing pores in the clay matrix rather than creating new ones. This finding is consistent with the sintering behaviour of silica-rich minerals described in section 3.2.5.

These results effectively disqualify diatomite as a viable pore-forming additive at 1100°C and directly inform the additive selection for the hybrid mix designs discussed in section 3.5.3.

Porosity by Diatomite Concentration

Diatomite (% wt)	900 °C	1000 °C	1100 °C
0%	27,44%	25,12%	16,60%
5%	29,19%	25,18%	13,73%
10%	29,59%	26,47%	14,37%
15%	32,71%	25,26%	11,59%

Table 3.5 — Apparent porosity by diatomite concentration.

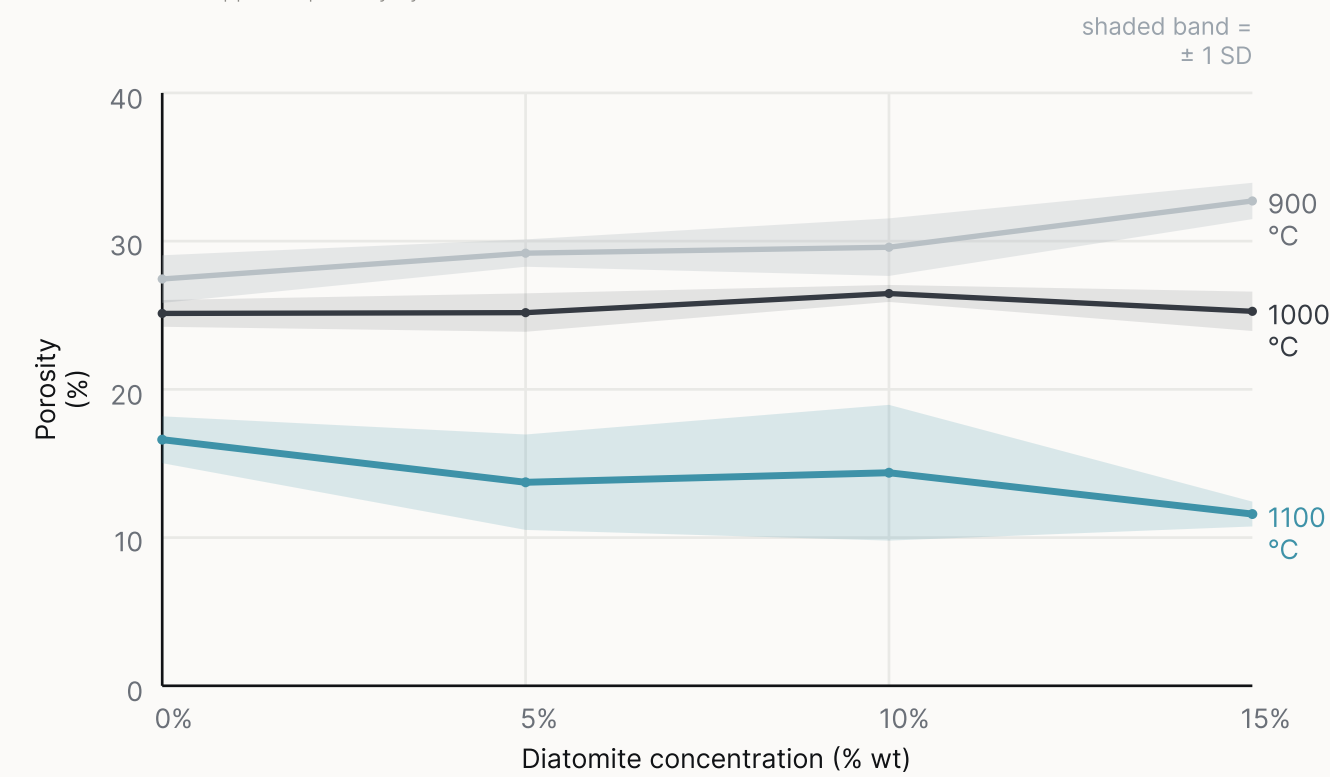


Fig. 3.13 — Apparent porosity vs diatomite concentration.

3.5.2 OPTIMIZATION FOR BIO-RECEPTIVITY AND FROST DURABILITY

For a dual-porosity ceramic to meet both the bioreceptive window (25–35% total porosity) and the frost criterion ($F_c \geq 70$), the total pore volume PV must divide into SCG-derived macropores ($V_{\text{macro}} = 0.291 \cdot PV$) and capillary micropores ($V_{\text{micro}} = 0.709 \cdot PV$). The clay already supplies part of that microporosity, so the volume left for a secondary micro-additive is the deciding quantity:

$$V_{\text{additive}} = 0.709 \cdot PV - V_{\text{micro_clay}}$$

If this leftover is negative anywhere in the target window, the clay's own micropores already use up the budget and no secondary additive can be added without pushing total porosity past the 35% ceiling.

Figure 3.14 plots V_{additive} against total porosity for the three firing temperatures. At 900 °C and 1000 °C the leftover is negative across the entire window and only turns positive at 37,4% and 36,7%, both beyond the ceiling, so both temperatures are eliminated. At 1100 °C, stronger sintering lowers the clay's baseline microporosity enough that the leftover stays positive across the full 25–35% window. 1100 °C is therefore the only temperature at which SCG and corn starch can be co-optimised within both constraints, answering RQ3. Because diatomite vitrifies at 1100 °C (Section 3.5.1), it cannot act as the secondary additive there, so the SCG–diatomite route is unachievable, answering RQ5.




Fig. 3.14 — Room for a secondary micro-additive.

3.5.3 FINAL RECIPE DERIVATION

The characterization data collected in sections 5.1.1 and 5.1.2, combined with the $F_c = 70$ framework established in section 3.2.6, provide all necessary inputs to derive an optimized ceramic recipe. The target total porosity was fixed at 30%, which represents the midpoint of the bioreceptive range. It also balances moisture retention against structural integrity. Based on this target, the required contribution of each component can be calculated as follows:

Applying the Maage relationship at $PV = 30\%$, a macropore fraction of 8.73% ($0.291 \times 30\%$) must be generated by SCG to satisfy the frost criterion. The remaining 21.27% constitutes the total required micropore volume, of which 16.60% is already provided by the baseline clay matrix firing at 1100 °C. This leaves a residual micropore contribution of 4.67% to be delivered by the secondary corn starch additive.

These target pore volume contributions were then mapped back onto the individual additive characterization curves at 1100 °C. Through linear interpolation between the measured data points, an SCG concentration of 5.66 wt% was identified as producing the 8.73% macropore volume required to satisfy the frost criterion. Similarly, a corn starch concentration of 3.95 wt% was found to generate the remaining 4.67% of additive-induced microporosity needed to reach the 30% total porosity target.

Prior to physical production, the theoretical recipe was refined for fabrication requirements. The white earthenware clay used in this study inherently contains approximately 25% chamotte. To improve dimensional stability during drying, reduce shrinkage cracking, and ensure consistent extrusion behaviour during 3D printing, the chamotte content was increased to 40%.

Figure 3.15 presents the effect of chamotte content on apparent porosity at 1100 °C. The results confirm that chamotte contributes only a marginal increase in porosity, rising from 16.5% at 25% grog content to 19.1% at 40%, representing a total increase of approximately 2.6 percentage points across the full tested range.

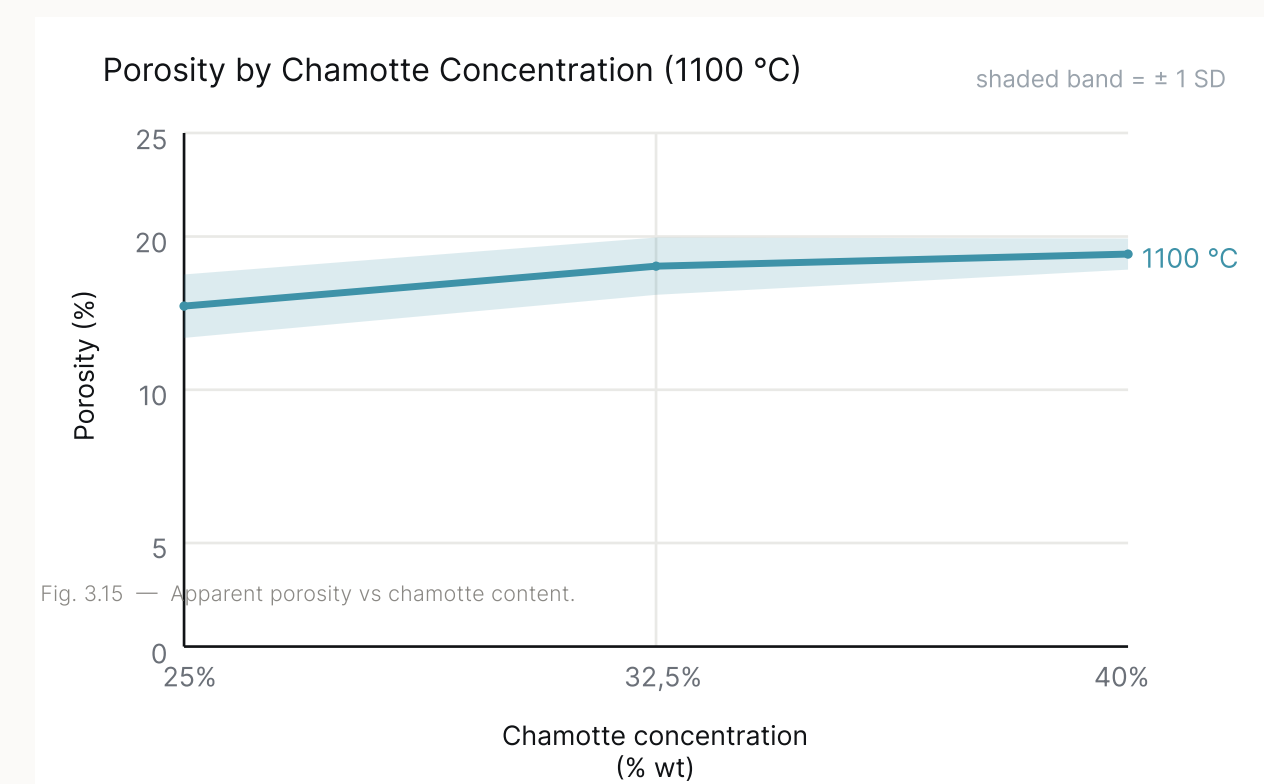


Fig. 3.15 — Apparent porosity vs chamotte content.

The resulting recipe is presented in Table 3.6. This composition simultaneously satisfies the bioreceptive porosity target of 25–35%, is calculated to meet the $F_c = 70$ frost threshold on an assumed bimodal pore structure, and operates within the only firing temperature at which both additives retain their pore-forming effectiveness.

It should be noted that the recipe derivation assumes the porosity contributions of SCG and corn starch are fully independent and additive. In practice, the two additives may interact within the clay matrix during firing, competing for the same void space, influencing each other's burnout behaviour, or jointly affecting the sintering kinetics of the surrounding ceramic. This could result in a final measured porosity that deviates slightly above or below the theoretical 30% target. The recipe should therefore be understood as a well-founded starting point rather than an exact prediction, and empirical validation through physical testing of the combined mix is presented in section 3.5.4.

Final Recipe		
Component	Concentration	Function
White earthenware clay	Balance	Structural matrix & baseline microporosity
Chamotte	40%	Fabrication stability; inert filler
Spent coffee grounds	5,66%	Primary macropore former (> 3,0 μm)
Corn starch	3,95%	Secondary micropore former (< 3,0 μm)
Firing temperature	1100 °C	Only viable sintering condition

Table 3.6 — Final recipe.




Fig. 3.16 — The fired ceramic recipe, 50x magnified.

3.5.4 FINAL RECIPE VALIDATION

POROSITY

The three samples of the final recipe returned apparent porosities of 41.56%, 37.33%, and 38.73%, giving an average of 39.21%. This exceeds both the theoretical 30% target derived in section 3.5.3 and the upper bound of the 25–35% bioreceptive range. The result confirms what section 3.5.3 anticipated as a risk. The three additives do not contribute independently. The combined interaction of SCG, corn starch, and the increased chamotte content of 40% during firing produces more porosity than the linear interpolation from the individual characterization curves predicted. While chamotte alone contributed only a marginal 2.6 percentage point increase in porosity across its tested range, its presence in the combined mix alongside two burnout additives likely influences the sintering kinetics and pore connectivity in ways the individual characterization experiments could not capture.

While this means the recipe derivation in section 3.5.3 should be understood as a starting point rather than an exact prediction, the result is not a failure. A porosity of 39.21% sits above the 25–35% target rather than inside it, but is still a highly porous ceramic capable of significant water uptake and retention. Further iteration on additive concentrations, specifically reducing either the SCG or corn starch dosage, would be needed to bring the result closer to 30% if tighter control over the pore volume is required.

CHEMICAL COMPATIBILITY

All ceramic samples made from the resulting recipe registered a pH of 6 across all time intervals (5 minutes, 1 hour, and 24 hours), identical to the demineralised water baseline. No chemical leaching was detected regardless of exposure duration or sample composition.

3.5.5 LIMITATIONS

Several limitations should be noted in the interpretation of these results. Sample fabrication started with LDM 3D printing for the initial SCG batches but was switched to hand-forming for all subsequent samples to allow faster iteration. This means the earliest SCG results were produced under different forming conditions than the rest of the dataset, which could affect pore distribution and compaction density in those samples.

Water-to-clay ratios were adjusted manually during mixing to reach a consistent extrudable consistency rather than measured precisely. This introduces a variable across batches that could independently affect the resulting porosity.

The Maage Factor calculation relies on a binary classification of pore sizes. This resulted in all SCG-derived porosity assigned to the macropore fraction and all starch and diatomite to the micropore fraction. In practice some starch pores will exceed the 3.0 μm threshold and some SCG pores may fall within it. This simplification, made necessary by the absence of mercury intrusion porosimetry data, likely underestimates the true F_c value and introduces a conservative margin into the frost resistance prediction.

Each sample condition was tested with a minimum of three samples, providing a reliable basis for averaging across conditions. However, without a larger sample size the results cannot fully account for natural variation in firing or mixing between batches.

04

PHASE 2 MESO-SCALE GEOMETRY



Fig. 4.1 — Experimental Meso-Scale Surface Topologies:
Tree Bark, False Diffusion, Spherical Undulation, Woven
Fabric, and Fractal Branching

Introduction

Phase One engineered a ceramic body that absorbs and holds water while staying chemically safe for colonisation and designed to resist frost. That settles how much water the material can store, but not whether the stored water ever reaches the organisms on the surface. In the water logic that structures this project, the micro-scale governs absorption and the meso-scale governs retention and availability, and it is the surface geometry, rather than the material, that decides how water is caught at the face, how long it lingers near the surface, and how well it is sheltered from drying. This phase therefore asks a single question: which printed surface geometry shows the most promise for keeping moisture available where the moss lives? The chapter first sets out the background and the design space the candidate geometries are drawn from, then develops five distinct geometries from the identical Phase One recipe, and finally tests them against one another for how they capture and retain water.

Background

Storage is not the same as availability. A highly porous body can draw water into its deeper structure, away from the surface where a rootless organism lives, so absorbed water is not automatically usable water; to be bioreceptive, a material must hold moisture at the surface long enough to stay accessible (Stohl et al., 2025). This "time of wetness" is the primary limiting factor for moss, which has no roots and depends entirely on its environment for hydration (Lim & Lharchi, 2023; Jakubovskis, 2024). The evidence that geometry, not bulk porosity, decides this is direct: in a comparative panel study, a flat surface made from highly porous, absorbent material showed almost no colonisation, while textured panels on the same material colonised readily (Mustafa et al., 2021), and textured concrete reached 60% moss coverage against 15% on a smooth equivalent over eight weeks (Moss Concrete, 2025). The meso-scale problem is therefore not how much water the tile holds, but whether its surface catches that water, holds it, and keeps it available where the moss can use it.

A surface manages this through its form, and the mechanism is well understood. Relief catches rain and slows its descent, while grooves, pits and crevices create sheltered microniches that protect spores and cells from the wind, gravity and runoff that would otherwise strip them away, easing their transition from loose to anchored (Bone et al., 2022; Parveen & Tariq, 2023). Roughness also increases the true surface area, which raises the capacity to trap rain, dew, dust and nutrients and to shade that moisture from evaporation (Mustafa et al., 2021; Rotondi et al., 2024). Crucially, the texture has to match the scale of the organism to work: studies of colonisation show roughness drives settlement only when its scale corresponds to the morphology of the colonising species (Bone et al., 2022), which links this phase directly to the organism-scale pore and groove targets established in Phase One. Two metrics quantify this capacity. Rugosity, the ratio of true to projected surface area, is 1.0 for a flat plane and rises with textured or printed relief; higher values mean more ecological space and a wider range of microconditions for colonisation (The Science of Bioreceptive Substrates, n.d.). Sz, the peak-to-valley height, captures the depth of shelter, and Mustafa et al. (2021) place its useful ceiling around 20 mm, beyond which deeper relief stops helping biologically and starts weakening the tile.

Additive manufacturing is what makes these features controllable. Casting fixes a surface in a mould; layer-by-layer extrusion lets the same process set the macro-form, the surface texture, and the internal porosity in one continuous operation (Bose et al., 2024). It also turns a defect into a tool: the horizontal layer lines left by printing act as continuous micro-grooves that retain water, gather debris, and shade the organisms inside them, and deeper grooves correlate directly with higher moss coverage (Lim & Lharchi, 2023; Soza Ruiz et al., 2025). The geometry of the print is therefore not just the shape of the tile but part of its water-handling mechanism.

Where the design space gets difficult is that theory offers several distinct geometric logics for doing this, with no clear winner for this task. Biomimetic surfaces abstract crevices and textures from bark, stone or the fog-harvesting Namib beetle to trap moisture passively (Sochůrková et al., 2023). Reaction-diffusion patterns such as the Gierer-Meinhardt model generate complex multi-scalar topologies tuned to organism scale (Abdallah & Estévez, 2023). L-system branching produces recursive channels that distribute water while keeping the wet clay stable during printing (Crawford et al., 2022). Sinusoidal undulations create alternating pockets and ridges that redirect water into deeper grooves and resist the rapid drying seen on flat surfaces (Mustafa et al., 2021).

4.2



Each logic is plausible, each handles water differently, by channeling, pooling or spreading it, and none has been tested head-to-head on a fired bioreceptive ceramic.

That is why this phase is a comparison rather than an optimisation. Rather than refining one form toward a calculated ideal, five geometries are drawn from these competing logics, fabricated from the identical Phase One recipe, and tested directly against one another for how they capture and retain water at the surface. The objective is to identify which geometric strategy best keeps moisture available where the moss lives, and the design space above defines the candidates that the next section develops into printable tiles.

Fig. 4.2 — Liquid Deposition Modeling: ceramic paste extruded layer by layer through the nozzle to form the printed tile.

4.3

Geometric Design and Fabrication

4.3.1 DESIGN APPROACH

The five geometries were each developed as a candidate drawn from one of the geometric logics set out in the background: a biomimetic bark texture, a reaction-diffusion pattern, two sinusoidal undulations at different frequencies, and an L-system branching structure. Translating these into printable tiles used a parametric pipeline in which a Python script generated the G-code toolpath directly, giving full control over how each form was deposited. Every design was previewed in a slicer and adjusted before printing. The scripting was collaborative, developed partly with AI coding tools and refined by hand, which made it possible to move from a geometric concept to print behaviour without authoring G-code manually.

After firing, all five tiles were 3D scanned to characterise the surface that was actually produced rather than the one that was scripted. Rugosity and Sz were extracted from the scan files, by Python analysis and direct measurement respectively, so that the comparison rests on the real fired geometry. This distinction matters because printing, drying and firing all deform the surface, so the physical tile never matches the toolpath exactly. One tile, G05 Fractal Branching, had debris lodged in its channels from earlier testing that caused the scan to under-read its depth; its rugosity and Sz were therefore measured manually from the 3D file using the highest and lowest identifiable surface points.

Not every geometry reached the comparison. Two genuine reaction-diffusion mesh topologies were attempted but produced severe extrusion inconsistencies during printing, the mesh complexity causing material build-up and surface artefacts that made the results incomparable. The five geometries presented here are the viable outcomes of that exploration, each built around a distinct logic for managing surface water and supporting growth. The reaction-diffusion logic still appears in the set, but as a printable Z-directional approximation (G02) rather than a true mesh, a substitution explained in its description below.

4.3.2 GEOMETRY MATRIX

The five geometries and their measured surface properties are summarised below. Each is described by its parent geometric logic, its surface character, the water mechanism it is designed to exploit, and the two measured metrics introduced in the background: Sz (shelter depth, with roughly 20 mm as the useful ceiling) and rugosity (true-to-projected surface area).

Tile Geometry Comparison

	G01 Tree Bark	G02 False Diffusion	G03 Spherical	G04 Woven Fabric	G05 Fractal Branching
Parent logic	Biomimicry (bark)	Reaction-diffusion	Sinusoidal (biaxial)	Sinusoidal (high f.)	L-system branching
Surface profile	Linear ridges	Organic bumps	Rounded lobes	Tight wave lattice	Branching channels
Feature scale	Medium	Medium	Coarse	Fine	Medium-coarse
Directionality	Vertical	Omnidirectional	Omnidirectional	Horizontal dom.	Vertical + branch
Water mechanism	Channeling	Pocket trapping	Deep pooling	Even spreading	Channel + pool
Sz (\approx 20 mm)	22,38 mm	16,15 mm	33,18 mm	18,43 mm	19,10 mm
Rugosity	3.6	4.0	4.6	4.2	4.3

Table 4.1 — Tile geometry comparison.

Two values sit outside the recommended depth window and are worth noting before any testing. G03 Spherical, at 33.18 mm, is well beyond the 20 mm ceiling Mustafa et al. (2021) identify, so its deep pockets may exceed the point of useful shelter and begin to compromise the tile. G01 Tree Bark, at 22.38 mm, sits just above the ceiling. The remaining three fall within or just under it. These are design properties of the fired tiles, not performance results; how they translate into water behaviour is the subject of the testing.

4.3.3 THE FIVE GEOMETRIES

Each geometry is described below by its design logic and the water behaviour it is intended to produce. These are statements of design intent; the tests in Section 5 examine how each tile actually behaved.

Description of design variations

G01 — TREE BARK



Fig. 4.3 — G01 Tree Bark.

A biomimetic translation of the vertical ridge structure of natural bark. Parallel channels run continuously up the tile, forming deep linear grooves between raised ridges. The intent is directional: the channels steer water downward in controlled streams while the ridges shade their own edges and offer continuous anchorage for rhizoids. Its regular, uniform texture also makes it the most predictable of the five, which suits it as a baseline against which the less ordered geometries can be read

G02 — FALSE DIFFUSION



Fig. 4.4 — G02 False Diffusion.

A printable stand-in for a reaction-diffusion pattern. A true reaction-diffusion mesh was attempted but caused extrusion failures, so the pattern was rebuilt by modulating the toolpath in the Z-axis to produce organic, rounded bumps rather than constructing it from a mesh. The result keeps the irregular, multi-directional pocket character of a diffusion pattern while remaining printable. The bumps form shaded cavities between them intended to trap moisture and debris without steering flow in any particular direction.

G03 — SPHERICAL UNDULATION




Fig. 4.5 — G03 Spherical undulation.

A sinusoidal waveform applied simultaneously in the X and Y directions, producing a field of large rounded lobes with deep concave pockets where the waves intersect. These pockets are the design's primary retention feature and are the deepest of the five geometries, intended to create strong localised shade and the largest enclosed volume for moss cushions to establish. As noted in the matrix, that depth also pushes it well past the recommended shelter ceiling, so its aggressive pooling is a deliberate extreme within the set rather than a balanced design.

G04 — WOVEN FABRIC



Fig. 4.6 — G04 Woven Fabric.

The same sinusoidal logic as G03 but at a much higher frequency, compressing the waves into a tight interlocking lattice that reads as woven fabric. Where G03 concentrates water in a few deep pockets, G04 is intended to spread it across a dense network of small niches, trading depth for distribution and presenting a large textured surface area relative to its footprint.

G05 — FRACTAL BRANCHING



Fig. 4.7 — G05 Fractal Branching.

A recursive L-system branching logic applied to the surface toolpath, producing vertical channels that split and rejoin as they travel up the tile. The branches widen and narrow at each bifurcation, creating alternating zones of concentrated flow and pooling, with junctions intended to act as collection points for moisture and debris. Its vertical orientation favours gravity-fed distribution, and of the five it most directly mirrors the macro-scale water network developed in Phase Three.

4.4

Methodology

This section sets out what the experiment varied, how the tiles were made and fired, and the three tests used to measure their water behaviour. Because the geometries are compared as single specimens, the section also states the limits of that comparison before any results are read.

4.4.1 RATIONALE VARIABLE SELECTION

Geometry is the only variable in Phase Two. All five tiles were printed from the identical Phase One recipe and fired on the same schedule, so material composition, porosity and pH are held constant and any difference in water behaviour can be attributed to surface form rather than to the material.

The tests are built around the three stages of the water cycle a facade tile has to manage, so that each test isolates one stage. The first is capillary uptake: how effectively a tile draws water upward from a reservoir below, which determines how well it sustains the moss from stored water between rain events. The second is moisture retention: how slowly a saturated tile releases water under summer heat and airflow, which governs how long the moss stays in its active, evaporating state after the store is filled. The third is surface capture: how much water a tile retains when rain strikes its face directly rather than running off, which determines how effectively a rain event refills the store. Together the three trace water through the tile across the full cycle, from rainfall, through storage, to availability at the surface.

4.4.2 EXPERIMENT DESIGN

Three tests were run across the five geometries, one for each stage above.

Experiment Design Table

Test	Property measured	Stage in the water cycle
Sorptivity	Rate of capillary uptake against gravity	Drawing on stored water between rains
Drying dynamics	Time of wetness under summer facade conditions	Holding the moss active after wetting
Rain simulation	Surface capture and retention from direct rainfall	Refilling the store during a rain event

Table 4.2 — Experiment design.

Because the five geometries differ in physical size and three-dimensional surface area, each tile's surface area was measured individually and all mass data was normalised per cm² to allow direct comparison.

One limitation is built into the design and shapes how the results must be read: each geometry was tested as a single tile. There is therefore no within-geometry replication and no statistical spread, so the results describe the behaviour of individual specimens rather than reliable averages. Differences between firing batches, mixing and water content, and print quality between tiles cannot be separated from the effect of geometry itself. The tests are accordingly treated as an exploratory, indicative comparison, and the results are interpreted as trends rather than established differences.

4.4.3 FABRICATION AND SAMPLE PREPARATION

All tiles were printed from the final Phase One recipe: white earthenware clay with 40% chamotte, 5.66 wt% spent coffee grounds, and 3.95 wt% corn starch. Each batch was mixed by hand until homogeneous, with water added until the paste reached a consistent extrudable stiffness.

The tiles were printed by Liquid Deposition Modeling using a 6 mm nozzle at a layer height of roughly 2.6 to 2.9 mm, settings drawn from the bioreceptive-printing literature as the range that maintains stable layer adhesion while preserving functional horizontal grooves (Sochůrková et al., 2023; Lim & Lharchi, 2023). These parameters were held identical across all five geometries so that the layer-line texture is common to every tile and only the designed macro-geometry differs.

4.4.4 FIRING PROTOCOLS

All tiles were fired in an electric kiln at 1100 °C with a one-hour soak at peak temperature, the same schedule used in Phase One. Keeping the firing protocol identical across both phases ensures that differences in water behaviour between the tiles come from geometry rather than from variation in the fired material.

4.4.5 TESTING AND ANALYSIS PROTOCOLS



SORPTIVITY (CAPILLARY ACTION)

Each tile was wrapped in plastic film leaving only the bottom 5 mm exposed, then stood vertically in a sealed container with that exposed edge submerged, so water could enter only through the base and travel upward against gravity. Mass was recorded at 1, 5, 10, 30 and 60 minutes, then at 2 hours, 3 hours 45 minutes, and 21 hours 15 minutes. Sorptivity (S) was derived from the linear relationship between cumulative absorbed mass per unit area and the square root of elapsed time, $\Delta m / A = S \times \sqrt{t}$, where Δm is the mass of water absorbed (g), A the surface area (cm²), and t the elapsed time (min). One limitation of this protocol is that water enters through the base rather than the front face, so it characterises uptake through the body as a whole rather than the surface geometry's own contribution to capture; that contribution is isolated by the rain simulation instead.



DRYING DYNAMIC (MOISTURE RETENTION)

After re-saturation, each tile's back face and side edges were sealed with tape so that water could escape only through the front face, replicating a mounted facade panel. Tiles were placed in a climate chamber at 30 °C and 65% relative humidity, oriented perpendicular to the internal fan so the front face received airflow without being aimed directly at it. To even out the uneven airflow in the chamber, tiles were arranged in a staggered front-back pattern and their positions rotated at every measurement, so each tile spent equal time at each distance from the fan. Mass was recorded every 15 minutes for the first half hour, then every 30 minutes, then hourly, up to 8 hours 15 minutes, and water loss was expressed as a percentage of total absorbed water.



RAIN SIMULATION (SURFACE CAPTURE)

Each tile was mounted at 25° from horizontal, set with a protractor. A 50 ml volume of water was applied over 3 minutes from a pipette held 30 cm above the tile and moved evenly across the full front face. Runoff was collected beneath the tile and measured volumetrically, and the retained volume taken as the difference between the 50 ml applied and the runoff collected.

4.5

Results and Discussion

This section reports the three tests in turn, then reads them together. Because each geometry is a single specimen, the results are presented as indicative trends rather than established differences, in line with the limitation set out in 4.2.

4.5.1 SORPTIVITY

Over the full test, the Fractal Branching tile absorbed the most water and drew it in fastest, reaching 3.8 g/cm² after 21 hours 15 minutes. The others followed in a closer group: Woven Fabric at 3.2 g/cm², False Diffusion at 3.0, Spherical at 2.9, and Tree Bark lowest throughout at 2.5. The clearest feature is the early slope: the Fractal Branching specimen pulled water in noticeably faster from the first minutes and held that lead for the whole test, while the other four began closely grouped before Woven Fabric and False Diffusion gradually separated from Spherical and Tree Bark. All five curves followed the expected ISO 15148 form, steep at first and flattening as the pore network saturated.

As noted in the protocol, this test admits water through the base, so it reflects uptake through the ceramic body as a whole rather than the contribution of the surface geometry to capture. The differences here therefore describe how each fired tile behaved as a whole specimen, and cannot be read as the effect of surface form alone.

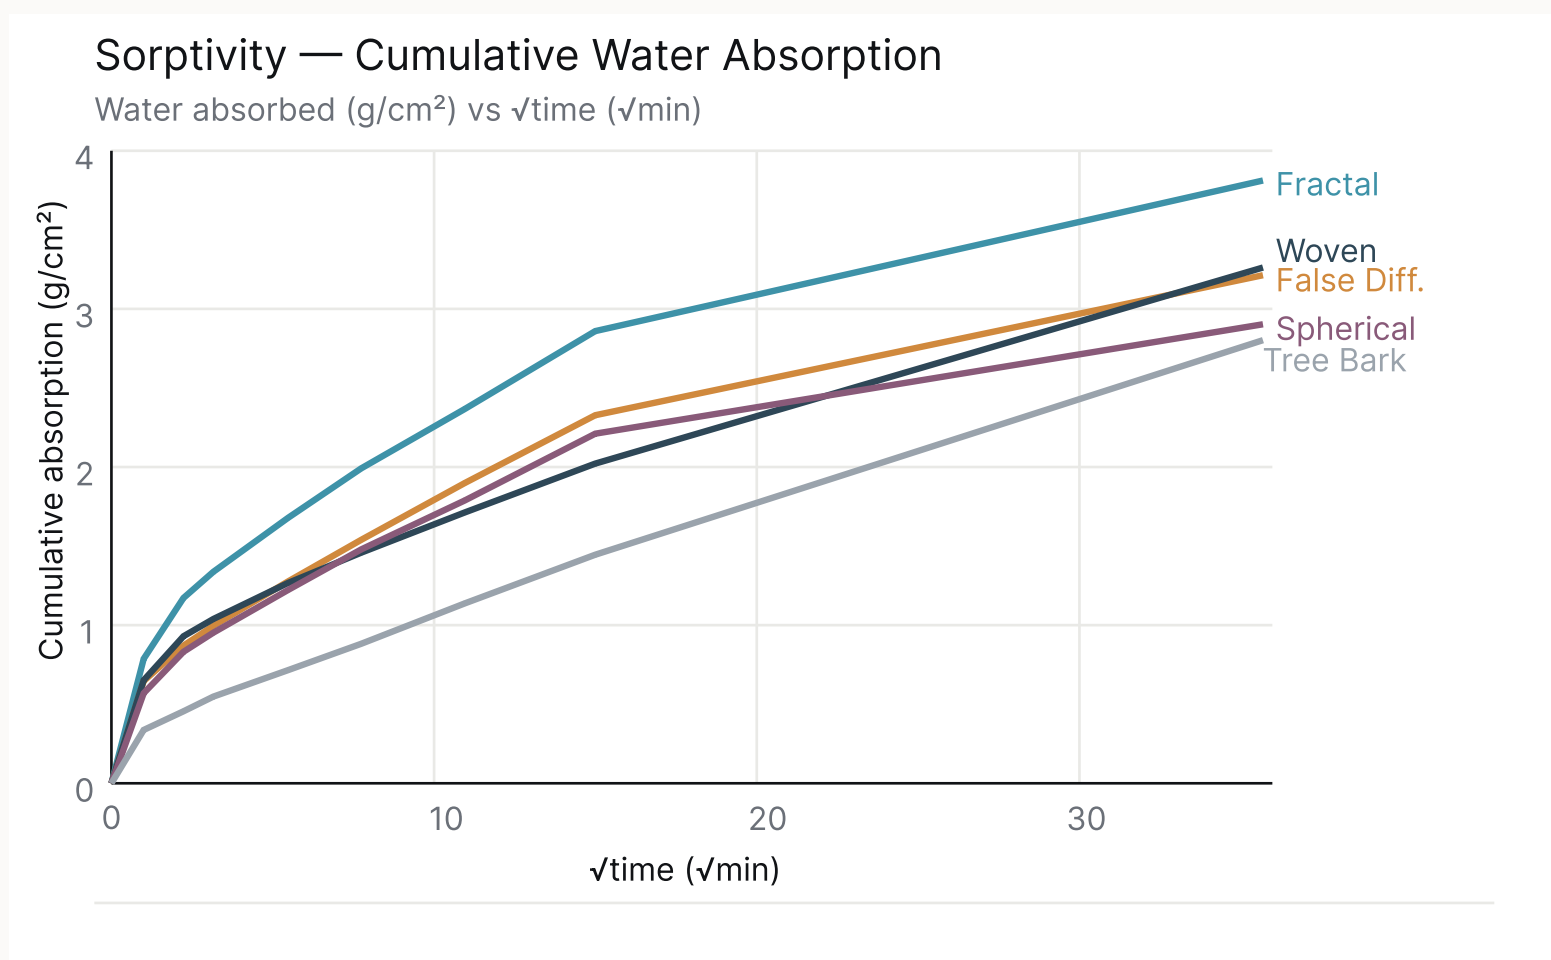


Fig. 4.8 — Sorptivity: cumulative water absorption.

4.5.2 DRYING DYNAMICS

With only the front face exposed at 30 °C and 65% relative humidity, the False Diffusion tile dried noticeably faster than the other four throughout, losing around 8.75% of its water content relative to dry mass by 8.25 hours. The remaining four stayed closely grouped, with Woven Fabric retaining the most at around 6.4% loss, Tree Bark at 6.75%, and Spherical and Fractal Branching both near 7.4%. Given that these four sit within roughly one percentage point of one another on single specimens, the differences between them should be treated as indistinguishable rather than as a ranking.

Some curves show a slight zig-zag between intervals. This is an artefact of the rotation protocol: tiles were moved between positions closer to and further from the fan at each measurement to even out the chamber's uneven airflow, so a tile temporarily lost water faster when more exposed and slowed when moved back. Some geometries appear more sensitive to this than others, which shows up as more pronounced variation in those lines.

The one result that stands clear of the grouping is False Diffusion's faster drying. Despite absorbing a reasonable amount in the sorptivity trial, this specimen released moisture from its front face more readily than the others under heat and airflow, which would work against sustained moss hydration on a mounted facade.

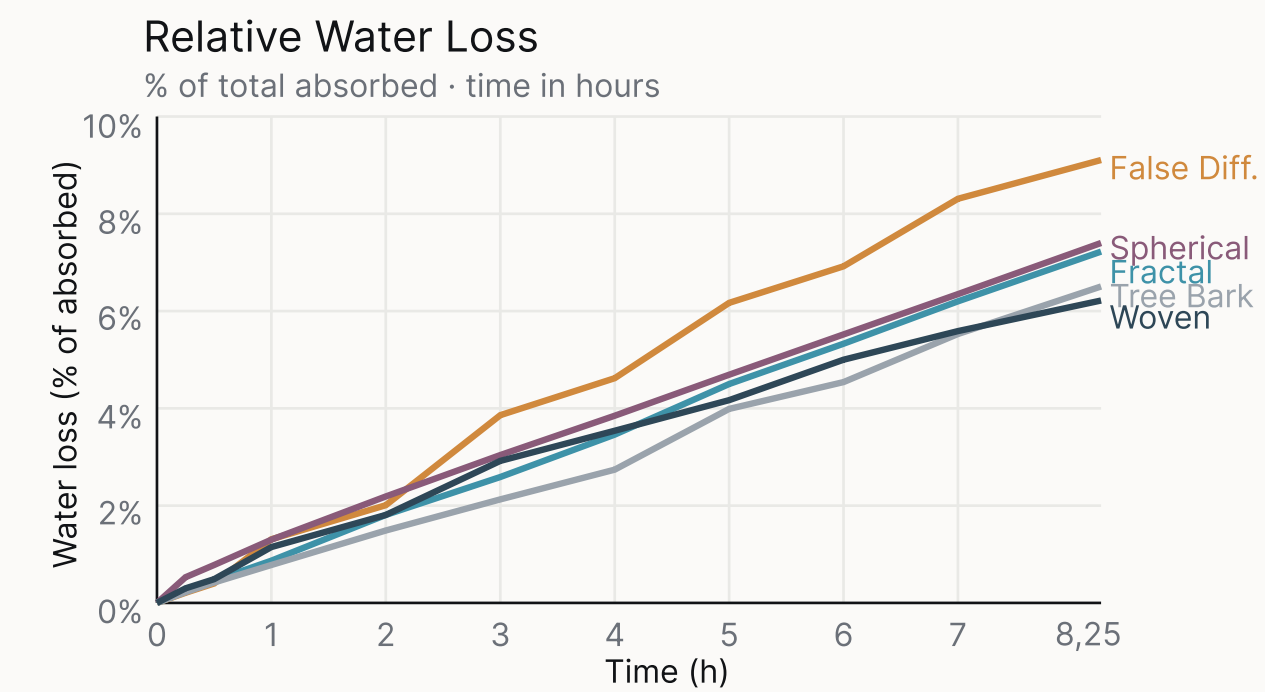


Fig. 4.9 — Relative water loss during drying.

4.5.3 RAIN SIMULATION

Three of the five tiles, Woven Fabric, False Diffusion and Fractal Branching, retained the full applied volume with no measurable runoff at the tested angle. The Spherical tile retained 89%, and the Tree Bark tile retained the least at 70%.

A consistent pattern sits behind these numbers. The three tiles that retained everything share an irregular, non-directional surface that places obstacles across any potential flow path, whereas Tree Bark's ordered vertical ridges offer little resistance to water draining straight down the channels. This is consistent with Mustafa et al. (2021), who found that surface irregularity increases the tortuosity of the flow path, slowing runoff and extending contact between water and substrate. The reading should stay at the level of this mechanism rather than the exact percentages, since with one tile each, the boundary between 89%, 100% and a repeat measurement cannot be established; what the test supports is that disordered, multi-directional surfaces resisted runoff better than the single directional one at this angle.

Water Retention

Geometry	Retention
G04 Woven Fabric	100%
G02 False Diffusion	100%
G05 Fractal Branching	100%
G03 Spherical	89%
G01 Tree Bark	70%

50 ml applied per sample
G03 runoff 5,5 ml · G01 runoff 15,2 ml

Fig. 4.10 — Rain simulation: water retention.

4.5.4 COMBINED READING

Read across the three tests, the Fractal Branching specimen showed the most favourable combination of behaviours: it absorbed the most and fastest in the sorptivity trial, retained the full applied volume in the rain simulation, and dried at a moderate rate comparable to the other high-retention tiles. Tree Bark sat at the low end of both uptake and rain retention, and False Diffusion, while mid-range on uptake, released moisture fastest in the drying test.

For a bioreceptive facade, the interval after rain stops is the one that matters, since during rainfall the moss is already wet and what decides survival is how much water the tile captured and how slowly it releases it. On the available evidence, the Fractal Branching tile performed favourably on both counts. It is important to be clear about the weight this conclusion can carry: each geometry was a single tile, the differences were not always large, and several fell within the margin that batch, mixing and print variation could plausibly produce. The result is therefore not that Fractal Branching is the optimal geometry, but that, among the specimens tested, it showed the most promising combination of uptake and retention. On that basis it was selected as the candidate to carry into the integration phase. Confirming whether its apparent advantage is real would require multiple samples per geometry, which is identified here as necessary further work.

This caution is not only about sample size but about the nature of the design space itself. The geometric possibilities for a printed surface are effectively limitless, since any parametric or biomimetic logic can be varied without end, and the five tiles tested here are a small purposive sample of that space rather than a representative survey of it. There is no single best geometry to be found, only forms that perform better or worse for a given organism under given conditions. The aim of this phase was never to identify an optimum but to compare distinct strategies and carry the most promising forward, and the result should be read in exactly those terms.

4.5.5 LIMITATIONS

The central limitation is sample size. Each geometry was represented by a single tile, so there is no statistical spread and the results describe individual specimens rather than reliable averages. Differences in firing, porosity distribution and print quality between tiles cannot be separated from the effect of geometry, and small differences between tiles cannot be distinguished from this background variation. Two test-specific limits also apply. The sorptivity test admits water through the base, so it does not isolate the surface geometry's own contribution to capture. And in the rain simulation, the three tiles that reached 100% retention may differ from one another in ways a single, fully-retaining trial at one angle cannot reveal; testing across a range of angles and rainfall intensities would be needed to separate them. Beyond sample size, the comparison is bounded by the design space it samples. Five geometries drawn from four design logics cannot represent the unbounded range of printable surfaces, and a geometry outside the tested set, or a variation on one within it, could plausibly outperform any tile here. The selection is therefore the most promising of the candidates examined, not a claim about the best geometry available. Across all three tests, the comparison should be read as an exploratory indication of how these geometries behave, providing a reasoned basis for selection rather than a statistically established ranking.

05

PHASE 3
MACRO-SCALE
SYSTEM INTEGRATION AND VALIDATION



5.1

Introduction

Phase Three brings the project together and tests it against its two goals: that the facade supports living moss, and that it cools. Phases One and Two could only measure proxies for these, porosity, sorptivity, drying rate, runoff, each a stand-in for a question that could not yet be asked directly. This phase replaces those proxies with the real outcomes, and completes the water logic that has structured the work: absorbed at the micro-scale, retained at the meso-scale, and now distributed across a facade.

This phase combines the validated recipe and the selected geometry into a single artefact and tests it against the conditions it was designed for. It is a synthesis-and-validation cycle in the Research-through-Design tradition, so its results are first functional outcomes and feasibility demonstrations, and its rigour comes from building every part of the system on constraints already established in the earlier phases.

The phase has two parts. The first, systemic integration, translates the single tile into a deployable module: an internal reservoir that buffers the moss through dry periods, continuous-column fabrication for seamless tiling, an interlocking assembly that distributes harvested water by gravity, and partial glazing that concentrates moisture where the moss lives. It is developed as a conceptual application, a demonstration of how the tile could become a system and of the constraints a production design would face. The second part, validation, proceeds in stages. Under real summer sun, a thermal comparison first establishes that the ceramic cools against a conventional wall, then weighs a glazed against an unglazed body for the balance each strikes between cooling and water retention; the moss is then introduced and compared on those surfaces. The chosen glazed module is afterwards taken into a climate chamber, moss-covered and bare, to measure how long it sustains moisture and cooling through a prolonged dry period. A biological study tests whether the substrate supports moss establishment. Together these close the two loops the introduction opened, turning cooling and biodiversity from claims into tested results.

5.2

Systemic Integration

This part translates the single test tile into a deployable facade module. It is developed as a conceptual application rather than a validated build: each feature is derived from a constraint already established in the earlier phases, and is presented as a demonstration of how the validated tile could become a system and of the requirements a production design would need to resolve.

The starting point is the geometry carried forward from Phase Two. Fractal Branching was selected there as the most promising of the tested geometries, with the caveat that this was an indicative result from single specimens rather than a proven optimum. It is taken forward here on that basis, both because it combined favourable uptake and retention and because its vertical, branching channel structure is the closest of the five to the gravity-fed water distribution a facade requires. A single test tile, however, is not yet a system. Four features are added to turn it into one.



Fig. 5.2 — Internal Cavity of the 3D-Printed Ceramic Tile Functioning as a Rainwater Storage and Buffering Basin

5.2.1 INTERNAL WATER RESERVOIR

The most consistent finding across the entire project, from the field research, the state-of-the-art review, and the Respyre case in particular, is that surface porosity alone cannot keep moss hydrated through an urban dry

period. The objective in Chapter 2 defined the system as self-sustaining only if it bridges that dry window without external watering. An internal reservoir is the feature that meets this requirement: a basin integrated into the module holds a volume of harvested water and feeds the porous body by

capillary action between rain events, buffering the moss precisely when surface moisture would otherwise run out. This is the single most important addition, because it is the one the proxies in Phases One and Two could not address and the one the validation study in Section 3 is designed to test.

5.2.2 SEAMLESS COLUMN GEOMETRY

Phase Two printed each geometry as a discrete tile, but a facade needs modules that tile continuously, so that the surface texture carries across module boundaries and the seams and layer-start points of printing do not become drainage or failure faults at the joints. The parametric G-code pipeline developed in Phase Two was therefore extended from single-tile generation to continuous multi-column module generation, producing the surface texture, the module envelope, and the internal reservoir geometry in one uninterrupted toolpath. This carries the same scripting method used throughout the project up to the module scale and demonstrates that the selected geometry can be fabricated as a module rather than only as a sample.

Fig. 5.3 — Continuous column geometry

keywords - assembly - mounting - glazing



Fig. 5.4 — Printed column module.

5.2.3 ASSEMBLY AND MOUNTING

A first assembly concept was developed to establish how modules mount to a facade and how water is managed across them. Each module carries notches along its sides so it can slide onto a system of rails, fixing the modules in a continuous vertical column. Each module is able to catch and store its own water, so at the unit level the tiles are self-sufficient. They are nonetheless linked into a column, with water entering at the top and flowing down through the stack by gravity, and the reason for this is maintenance rather than supply.

The driver is frost. A porous ceramic holding standing water is vulnerable to freeze-thaw damage in winter, when water expands as it freezes and can crack the body, the same risk the material's frost-targeting recipe was designed to mitigate at the pore scale in Phase One but which a full reservoir of standing water reintroduces at a larger scale. The system must therefore be taken out of service over winter and drained. If every module were fully independent, each would have to be closed and emptied individually. By connecting them into a top-fed column, a single closure at the top shuts off and drains the entire column at once, turning winterisation from a per-tile task into a single action.




Fig. 5.5 — Gravity-fed column assembly.

At this stage this is a proof of concept rather than an engineered connection. It shows a plausible way the modules could be mounted, watered and winterised as a column, and surfaces the constraints a production detail must resolve: load transfer and tolerance at the rail, sealing and water hand-off between modules, complete drainage for frost protection, and even distribution down the height rather than over-supplying the upper modules and starving the lower ones, the same gradient failure observed in the passive soil-based systems in Chapter 1.



Fig. 5.6 — Partial glazing detail, in earlier blue-glazed variant.

5.2.4 PARTIAL GLAZING

Glazing was held off entirely through Phase Two so that the geometry comparison isolated geometry alone. It is reintroduced here as a deployment refinement applied to the selected geometry. The treatment glazes the exposed peaks of the surface, which shed water and dry fastest, while leaving the recessed valleys unglazed and porous. This separates the shedding function from the retaining function, concentrating moisture in the porous valley zone where the organisms establish and where it stays available to rootless moss. It is a direct application of the second field principle, in which the boundary between glazed and porous areas follows the boundary between shedding and retaining geometry, and its effect on both water retention and moss hydration is one of the variables the validation study sets out to test.

5.2.5 PARAMETRIC GENERATION

This part describes the computational pipeline behind the facade. A Python script generates the tile geometry, produces the printer G-code, and builds an interactive preview where each tile can be inspected and downloaded for fabrication. Everything happens in one run, from pattern to machine file.

The surface pattern comes from a recursive branching algorithm. It is seeded at regular points along the full width of the wall, and a pair of sine and cosine functions shifts the branch angle and length on every layer, so the relief changes gradually as the wall rises. The script builds this pattern across the whole facade first and then cuts it into individual tiles. Because the geometry is continuous before it is divided, the branching lines carry straight across the joints and the tiles read as one surface. The functional features from the design are all parameters in the script: the wavy front face that catches and slows water, the corrugated back wall, the interlocking notches for assembly, and the drain hole.

The script writes the machine G-code itself. Each tile prints as one continuous spiral on a single solid base layer, with extrusion rate, feedrate and layer height set for a six millimetre ceramic-paste nozzle. Because the toolpath comes straight from the geometry, the print lines themselves form the water-catching grooves, and the bead stays under full control the whole way.

The same run assembles every tile into an interactive web viewer. The full six by six facade can be rotated and inspected, a tile selected, and its G-code downloaded in the browser. Raised ridges are shaded pale and the recesses moss green, so the places where water and growth would gather are clear at a glance.

The pipeline shows how the system could be deployed. It is an early demonstration, and a production version would need more work, but it already turns one validated tile into a facade that can be regenerated at any size, adjusted for its context, and sent straight to a printer.

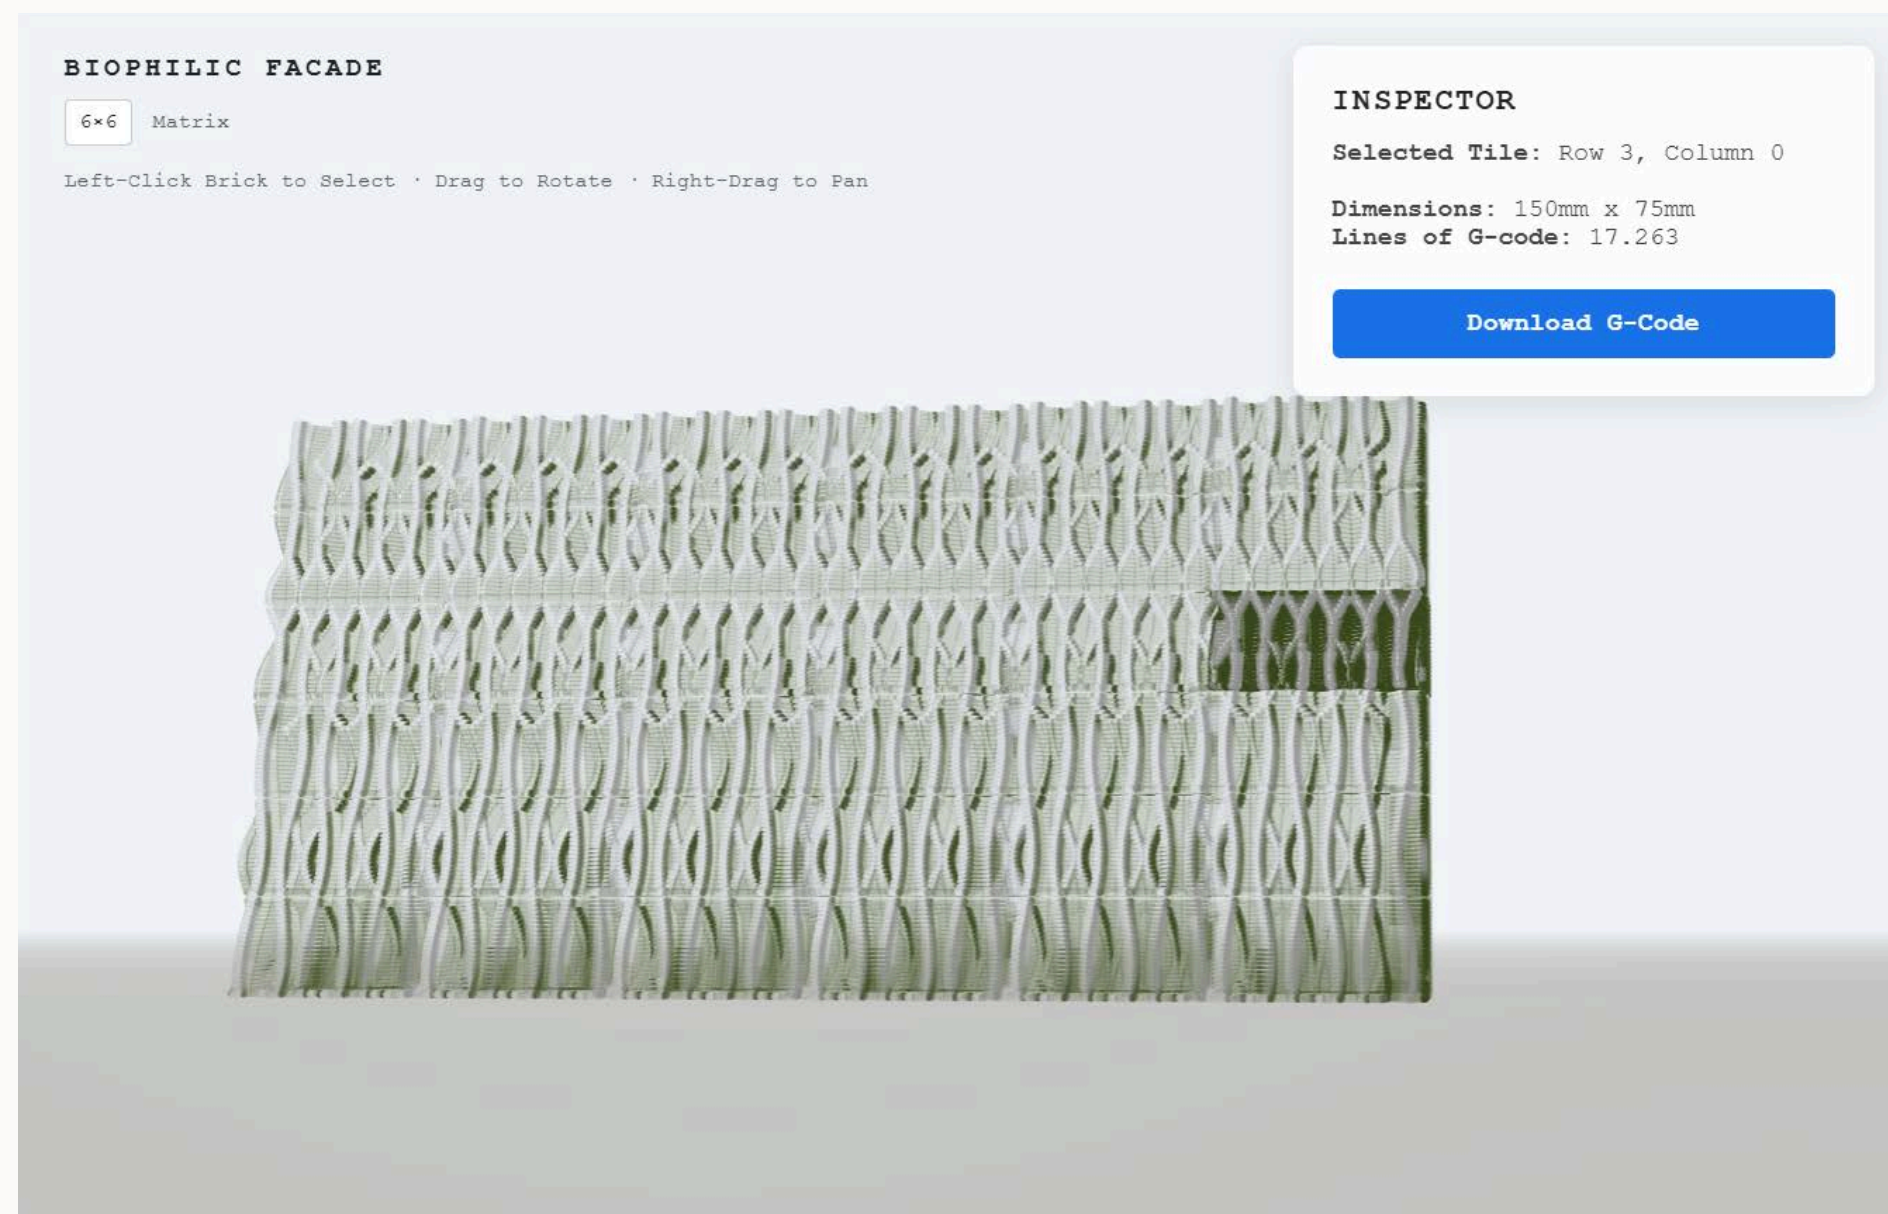


Fig. 5.7 — Interactive Three.js preview of the full 6 × 6 facade; any tile can be selected and its G-code downloaded in the browser.

5.3

Methodology

This section sets out how the integrated module was validated against its two goals, cooling and biological establishment. It explains what the studies varied, how they were staged, how the tiles were prepared and moss established, and the protocols used to measure each outcome. Because the prototype was tested as a small number of specimens under partly uncontrolled conditions rather than as a controlled factorial series, the limits of the comparison are stated alongside the results.

5.3.1 RATIONALE VARIABLE SELECTION

Two variables governed the validation: the surface treatment, glazed or unglazed, and the presence of the moss layer. Rather than crossing both at once, the studies were staged. The surface treatment was settled first, on bare specimens, by weighing the cooling each gave against the water it retained. The glazed treatment was then carried forward, and the moss became the only remaining variable, compared first outdoors for its cooling contribution and then, on the chosen glazed module, measured in a chamber for how long it could be sustained. Keeping the surface treatment fixed after the first stage meant each later comparison isolated a single variable.

5.3.2 EXPERIMENT DESIGN

The validation proceeded in three staged comparisons, supported by a biological establishment trial.

Stage	Specimens	Setting	Question
1	Glazed and unglazed stacks, plus a conventional wall (bare)	Outdoor, full sun	Does the ceramic cool, and which surface treatment to carry forward
2	Glazed-moss and glazed-no-moss stacks	Outdoor, full sun	The cooling contribution of the moss on the chosen surface
3	Glazed-moss stack	Climate chamber	Whether the module meets the self-sustaining target, and for how long

Table 5.2 — Experiment design.

Stage 1 established whether the ceramic cooled at all, by measuring it against a conventional sun-facing wall, and selected the surface treatment by comparing a glazed and an unglazed column placed side by side in the same stack, so that both met the sun under identical conditions and the glaze was the only difference between them. With glazing carried forward, Stage 2 introduced moss and compared a moss-covered glazed specimen against a bare glazed one under the same outdoor conditions, isolating the cooling contribution of the living layer. Stage 3 took the selected glazed-moss module into a controlled chamber and measured how long it sustained moisture and kept the moss active through a prolonged dry period, assessed against the dormancy threshold and the self-sustaining target defined in Chapter 2 rather than against a bare reference, since the moss contribution had already been established outdoors. The biological establishment trial ran in parallel, monitoring whether transplanted moss attached and survived on the substrate.

5.3.3 FABRICATION AND SAMPLE PREPARATION

All specimens were printed from the validated Phase One recipe, which had a measured apparent porosity of approximately 39%, and were fired on the same schedule used throughout the project. Tiles were assembled into columns of three, each with a filled reservoir, to approximate a section of wall.

For Stage 1, a single stack was built from two side-by-side columns, one glazed on the surface peaks and one left fully unglazed, so the surface treatment was the only difference between them and both columns sat under identical sun, orientation and position. The glaze was transparent, so the two columns looked almost alike and the difference between them was functional rather than visual. For Stages 2 and 3, glazed specimens were prepared in two states, one with moss established on it and one left bare, with the moss-covered specimen carried into the chamber for Stage 3.

Moss was gathered locally and attached to the porous valley zones, in the same coverage and positions across specimens. The attachment slurry was chosen through a short preliminary trial: an initial buttermilk-based slurry, intended to supply nutrients, produced extensive mould, an expected outcome given its high nutrient content, so the binder was changed to a cornstarch-based slurry, whose lower nutrient content markedly reduced mould while still holding the moss in place. This cornstarch slurry was used for all subsequent samples. After application the moss was left to establish for 5 days and kept continuously damp so it stayed active rather than drying into dormancy. Moss source, slurry composition and rest period were recorded.

5.3.4 TESTING AND ANALYSIS PROTOCOLS

OUTDOOR THERMAL COMPARISON

The stack was placed outdoors in full sun on a clear day of around 30 °C, alongside a conventional wall of the same orientation that was monitored across the day as a reference. Surface temperature was recorded by infrared thermography using a FLIR E6 PRO thermal imaging camera. These images show how temperature was distributed across the two columns and their peaks. Water loss was determined by measuring the water volume remaining in the reservoirs at the end of the day. This protocol served both the bare comparison in Stage 1 and the moss comparison in Stage 2.

CLIMATE CHAMBER LONGEVITY

For Stage 3 the glazed-moss module was dried to a constant mass and its reservoir filled to capacity, then placed in a climate chamber at 30 °C and 60% relative humidity. The module was weighted over the course of 20 hours to track the evaporation. The run continued until the drying trend was established, and the remaining duration was extrapolated from the curve rather than run to completion.

BIOLOGICAL ESTABLISHMENT

The transplanted moss was monitored across the establishment period for attachment, colour and overall health, and for the persistence of a living, hydratable layer. Survival was assessed visually, the moss judged viable where it stayed green and rehydrated readily after wetting.

5.4

Results and Discussion

This section reports the validation in the order the stages were run, from the bare thermal comparison through to the living module. The Stage 1 results are complete; the moss and longevity stages are reported as they stand



Fig. 5.8 — Cooling experiment set-up, glazed v.s. unglazed ceramic modules

5.4.1 SYSTEM AND CONVENTIONAL SURFACE

Under full sun at around 30 °C, both ceramic columns ran far cooler than the conventional wall of the same orientation. Across the day the wall climbed to roughly 50 °C, while the ceramic surfaces stayed near 28 °C and, importantly, held that temperature almost constant as the day went on rather than rising with the sun (Fig. 5.9). The evaporating surface buffered itself against the heat instead of tracking it, where the wall heated up through the day, the ceramic kept a stable temperature throughout.

In the infrared images this reads at a glance: the wall glows orange and yellow while the ceramic stack sits deep blue against the same warm scene. Before any biological layer is added, the porous, water-holding ceramic is already a dramatically cooler and more thermally stable surface than a conventional facade under identical sun, confirming the premise the whole system is built on.




Fig. 5.10 — Infra-red image of traditional concrete facade with glazed ceramic tiles




Fig. 5.9 — Cooling effect glazed v.s. unglazed modules

5.4.2 EFFECT OF SURFACE TREATMENT

Between the two columns the difference was small in temperature but large in water. The unglazed column was the cooler of the two at about 26.0 °C, against 27.5 °C for the glazed, a gap of 1.5 °C. After one day in the sun, however, the glazed column retained about 150 ml of water against 60 ml for the unglazed, more than twice as much. Because the glaze was transparent the two columns looked almost identical, so this difference is functional rather than visual, and it is best read as a trade-off rather than a verdict on cooling alone. The unglazed surface buys slightly lower temperature at the cost of losing water more than twice as fast, while the glazed surface gives up 1.5 °C to retain far more moisture and so sustain its cooling much longer. Since the system depends on bridging dry periods, the glazed treatment was carried forward, the small thermal cost being justified by the large gain in water retention.




Fig. 5.11 — Glazed ceramic module, temperature reading, infra-red




Fig. 5.12 — Unglazed ceramic module, temperature reading, infra-red




Fig. 5.13 — Glazed peaks, showing a relatively higher temperature than the surrounding ceramic




Fig. 5.14 — Unglazed peaks, showing a relatively lower temperature, due to a higher evaporation rate near the peaks

5.4.3 SURFACE EVAPORATION BEHAVIOR

The infrared close-ups make the mechanism behind this trade-off directly visible (Fig. 5.13 and 5.14). On the unglazed surface the peaks read as the coolest points, because water evaporated fastest there, which is also why the unglazed column lost water more quickly overall. On the glazed surface the pattern reversed, the sealed peaks shed and retained rather than evaporating, showing warmer in the same image. This is direct visual evidence for the partial-glazing logic: glazing the peaks suppresses the fastest evaporation while leaving the valleys porous for retention and moss, capturing both the cooling of the unglazed surface and the longevity of the glazed one.

5.4.4 BIOLOGICAL ESTABLISHMENT

Each step below records how living moss was established on the printed surface, from gathering it to growing it in. The method was developed by trial, so what failed is recorded alongside what worked, since arriving at a workable way to attach the moss was itself a result of this phase.

The establishment process

01 — COLLECTING THE MOSS



Fig. 5.15 — Moss growing on rough concrete surface with imbedded rocks.

Moss was gathered from hard, rocky and vertical surfaces around the site rather than from soft ground. The assumption was that species already living on stone and walls would adapt most readily to a hard, vertical ceramic facade, as they were already suited to thin substrates, anchorage by rhizoids and exposure, rather than to rooting in soil.

02 — BLENDING THE MOSS SLURRY



Fig. 5.16 — Blended moss slurry containing buttermilk and moss leaflets

An initial slurry was made with buttermilk, a common recipe intended to feed the moss, but it grew mould heavily and smelled, and was abandoned. The binder was switched to cornstarch and water, heated until it thickened into a gel and then left to cool, after which the moss, broken down into its leaves and fragments, was mixed through. The lower nutrient content kept the mould down while the gel still held the moss in place.

03 — APPLYING THE MOSS SLURRY



Fig. 5.17 — Moss slurry applied on fired ceramic

An earlier iteration split the surface to compare two approaches: one half coated with slurry alone, the other with whole moss fragments worked into the slurry. The slurry-only half never developed visible growth within the project's timeframe, since colonising from suspended spores and fragments is slow. This ruled out slurry on its own and confirmed that pieces of living moss had to be attached directly.

04 — ATTACHING THE MOSS



Fig. 5.18 — Moss attached on the ceramic tiles using the moss slurry as glue

Following that result, fragments of living moss were pressed directly into the slurry in the porous valleys of the tile, where moisture is retained and the surface is sheltered. Attaching established moss rather than waiting for colonisation gave the layer a living head start and concentrated it in the zones the geometry was designed to keep damp.

05 — WATERING & GROWING



Fig. 5.19 — Close-up internal water basin - ceramic module with established moss

Once attached, the moss was watered and misted regularly and kept continuously damp so it stayed active and could anchor and spread rather than drying into dormancy. It was first tested after five days of establishment and has since remained healthy for more than a month at the time of writing.

5.4.5 COOLING WITH THE LIVING LAYER



Fig. 5.19.5 — Cooling experiment set-up, glazed moss v.s. glazed bare ceramic modules

Read as a Research-through-Design probe rather than a controlled experiment, this stage was less about proving a cooling figure than about discovering how the integrated tile actually behaves in the sun. A single stack was built with two side-by-side columns, the left carrying established moss and the right left bare, and tracked with the infrared camera across a day alongside bare urban surfaces used as baselines (experiment set-up, left). Because it was a single prototype read with spot measurements under real, uncontrolled weather, the readings are treated as indicative observations that inform the design rather than as generalisable values.




Fig. 5.20 — Early evening, around 20:00: the bioreceptive tile sitting at roughly ambient temperature.




Fig. 5.21 — Early evening: a bare brick facade still reading 46.9 °C, far above the surrounding air.

The first observation was confirming: through the whole day the tile stayed far below the bare surfaces, and tracked the air rather than the sun. In the late morning it read in the high twenties while a bare wooden panel was already at 37–42 °C. At midday it sat in the low-to-mid thirties against bare decking at 48 °C. The contrast was sharpest in the evening: with the air at 26 °C around 20:00, the tile had cooled to 24–25 °C, at or just below ambient, while a bare brick facade still read 46.9 °C, more than twenty degrees above the surrounding air (Figure 5.20, tile; Figure 5.21, bare wall). By the next morning the tile was near 20 °C. The ceramic-and-water system therefore did not just stay cooler under peak sun, it tracked the air temperature and shed its heat after dark, where the bare surfaces went on releasing stored heat long after sunset.




Fig. 5.22 — Midday: a bare, conventional sun-facing surface reaching around 48 °C.




Fig. 5.23 — Midday: the moss column at around 35 °C.




Fig. 5.24 — Midday: the bare ceramic column at around 30.5 °C, cooler than the moss beside it.

The second observation was a surprise that, once paired with the water readings, reshaped how the moss is understood. While the sun was on it, the moss column ran consistently warmer than the bare-ceramic column: around 30 °C against 25.6 °C at the middle in the morning, and around 35 °C against 30.5 °C at midday (Figure 5.23, moss column; Figure 5.24, bare column; Figure 5.22, conventional surface). Even at its warmest, though, the moss column stayed roughly ten degrees below the conventional sun-facing surfaces measured at the same hours, so this was a difference only relative to the bare ceramic beside it, not an absolute heat penalty; the moss never approached the temperature of a normal wall. After sunset the difference vanished, both columns sitting near 25 °C in the evening and near 20 °C the next morning. The water, however, told the opposite story. By the following morning the moss column still held about 100 ml of water against only

about 20 ml for the bare column, roughly five times as much retained. The bare ceramic had cooled harder by evaporating its water away, while the moss column stayed warmer because it evaporated far less, and so conserved its store.

This locates the cooling in the wet ceramic and recasts the moss as a water-conserving layer. The bare ceramic exposes its pores to the air and evaporates freely, cooling hardest but draining its store close to empty within a day, while the moss caps that surface evaporation, giving up a few degrees of daytime cooling in exchange for holding several times the water. It is the same trade-off the glaze produced in Stage 1, slow the surface loss and the effect is sustained for longer, here carried out by the living layer rather than the coating. The moss also did not trap heat at night: it cooled to ambient like the bare column, most likely because .

the retained water kept it from drying into a dormant, insulating state

For the design argument this is a positive result. It places the cooling firmly in the ceramic body, and gives the moss two roles beyond urban biodiversity and fine-dust capture: it conserves the stored water, extending how long the system can keep cooling through a dry period, and it does so without the night-time heat-trapping that dry moss causes on conventional surfaces. The moss's daytime warmth is therefore not a liability but the visible counterpart of water being conserved rather than spent, and the longevity this implies is exactly what the chamber stage goes on to measure. Two questions remain for a further cycle: whether a fully hydrated, actively transpiring moss would cool more under peak sun, and how the balance between evaporation and conservation shifts as the store depletes.

Surface temperature across a summer day

Bioreceptive tile vs bare surfaces, 22–23 June, Delft - air from KNMI station 344 (Rotterdam)

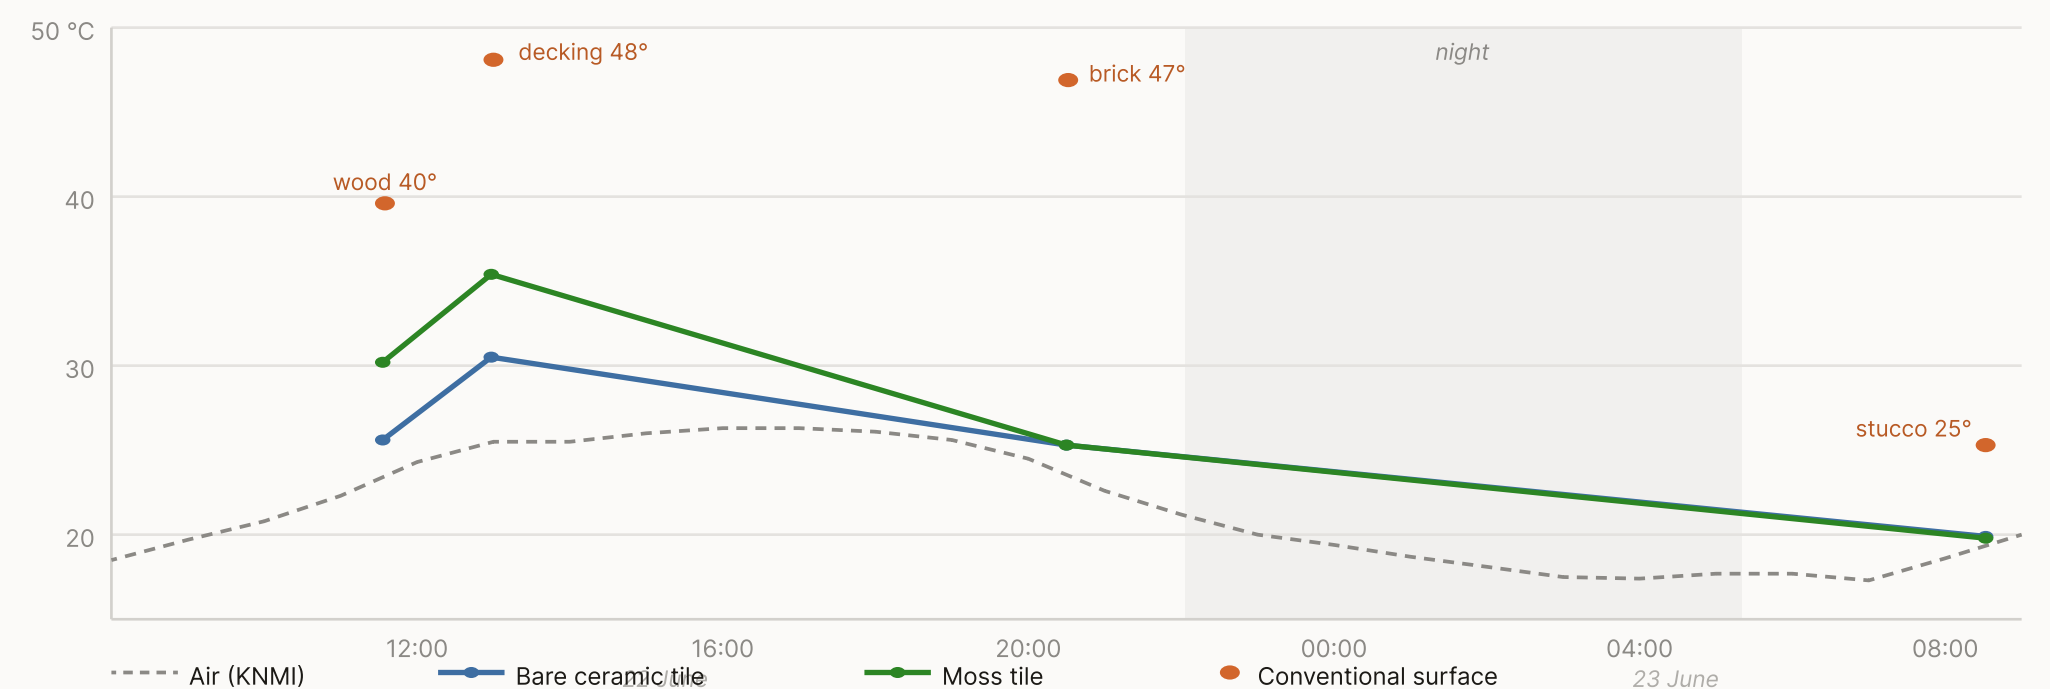


Fig. 5.25 — Graph showing the mentioned data points, in comparison to the ambient temperature

5.4.6 LONGEVITY OF THE LIVING MODULE



Fig. 5.26 — Climate-chamber set-up: the saturated glazed-moss module dried at a constant 30 °C and 60 % RH.

Where the outdoor stages tested the module against real, fluctuating weather, the chamber stage held it under a single relentless condition to see how long it could sustain itself unaided. The selected glazed-moss module was saturated and then dried at a constant 30 °C and 60 % relative humidity, weighed at intervals across nineteen hours (Fig. 5.25). Saturated, the porous body had taken up about 600 ml into its internal porosity. It then lost water steadily: roughly a quarter was gone within five hours, and about 79 % after nineteen hours, leaving some 125 ml still held in the pore network. Extrapolating the drying curve (Fig. 5.26), full depletion fell around the one-day mark.

Taken at face value this is short of the self-sustaining target set in Chapter 2, of bridging a multi-day peak-summer dry period. But the chamber condition is deliberately harsher than reality, and reading the two validation environments together is what makes the result useful. A constant 30 °C with continuous airflow and no night gives the surface no relief, where the outdoor cycle does: the cool, humid nights and the dew that forms on them slow evaporation sharply and can even return water to the surface. The Stage 2 module made this concrete, outdoors, the same kind of module still held about 100 ml the next morning, far more than the chamber would predict over the same span. The chamber figure is therefore best read as a conservative lower bound on endurance rather than a literal field lifetime; real performance sits above it.

What the test did establish clearly is the gap between how much water the module can hold and the dry periods it is meant to bridge. Even allowing that reality is kinder than the chamber, a porous body that gives up most of its stored water in around a day under continuous load is undersized for a multi-day drought. This points to two design responses. The first is to carry more water, through a larger or deeper porous body, or by adding a dedicated reservoir basin to the module rather than relying on the pore network alone. The second, and the more promising, is to stop depending on a single fill and add passive rainwater buffering, so the store is replenished from roof or facade runoff rather than only depleted, of the kind already demonstrated by the gravity-fed rainwater buffers in use at The Green Village (Fig. 5.27). That route folds the rainwater-harvesting logic seen in the field research and the state of the art back into the module, and shifts the self-sustaining question from how large a single store must be to how a continuously topped-up store could keep the moss active across a prolonged dry spell.



Fig. 5.28 — Rainwater buffer in use at The Green Village.

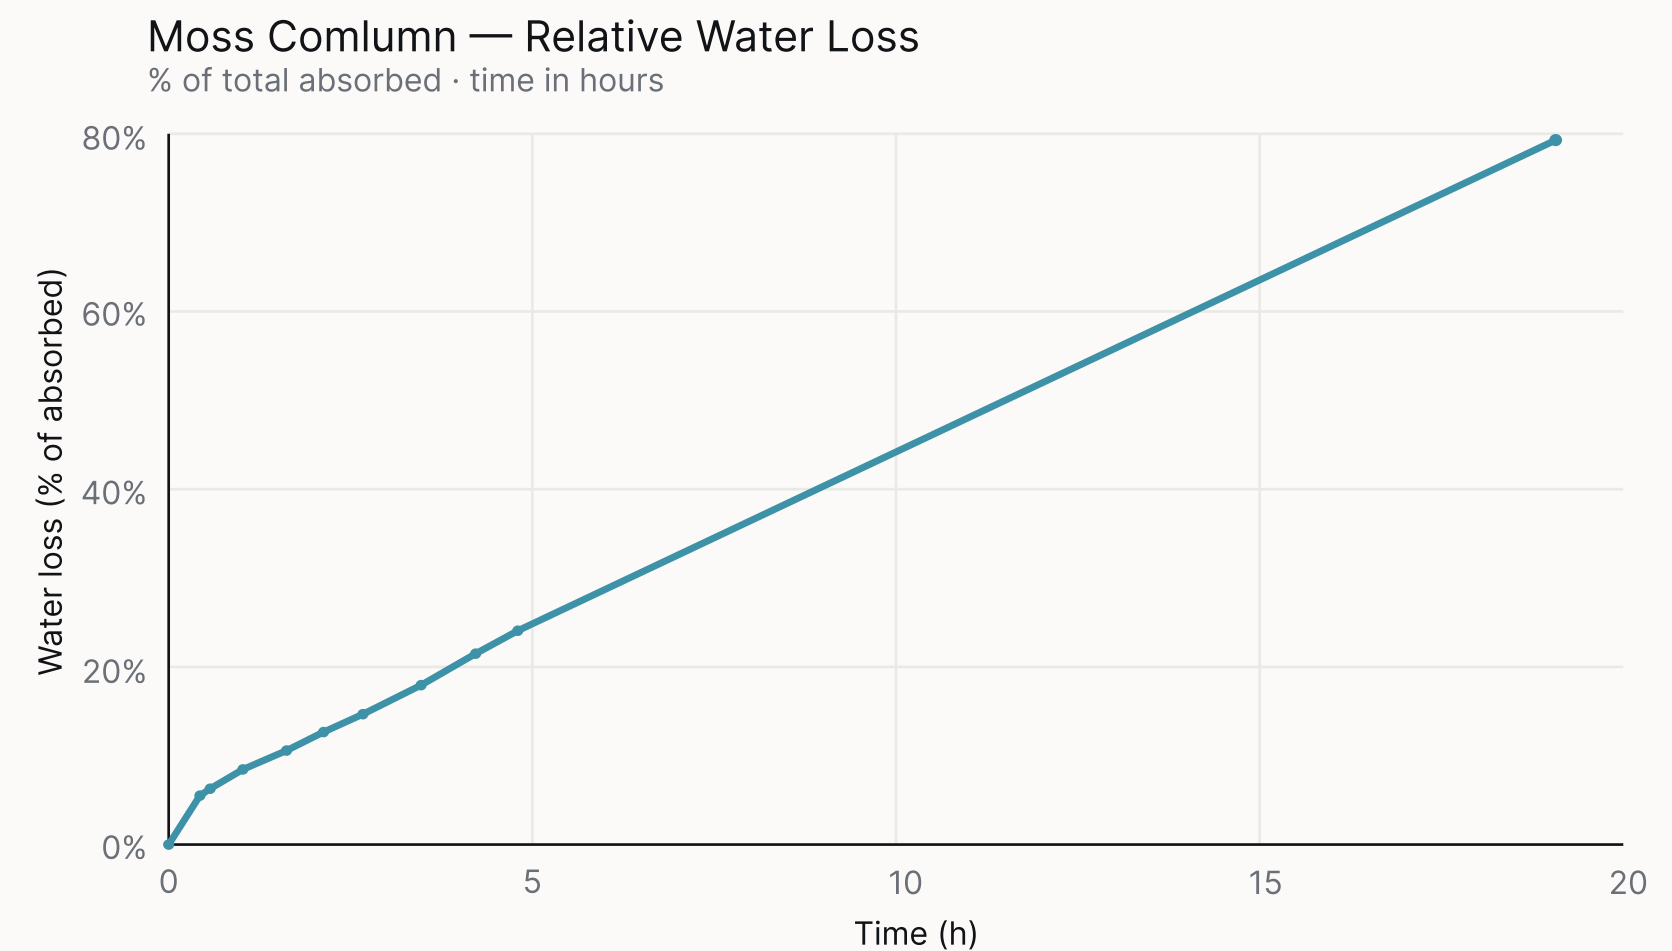


Fig. 5.27 — Relative water loss from the moss column over nineteen hours of constant-condition drying.

5.4.7 COMBINED READING

Across the three stages a consistent picture emerged, and it differs from the simple expectation the project began with. The cooling is not done by the moss but by the wet ceramic. Under full sun the ceramic surfaces sat more than twenty degrees below a conventional wall, held close to the air temperature rather than climbing with it, and shed their heat after dark instead of storing it, in the evening the tile sat at or just below the 26 °C ambient while a bare brick facade still read 47 °C. That is the heart of the result: a porous, water-holding ceramic is a fundamentally cooler and more thermally stable surface than a conventional facade, by day and by night.

What the glaze and the moss each contribute is not extra cooling but longevity, and they do it the same way. Both slow the loss of water from the surface, trading a small amount of immediate cooling for a much larger gain in retention. The glazed column ran about 1.5 °C warmer than the unglazed but held more than twice the water after a day in the sun, 150 ml against 60 ml. The moss column ran a few degrees warmer than the bare ceramic under sun, yet kept roughly five times the water by the next morning, 100 ml against 20 ml. In both cases the warmer surface is simply the visible sign of water being conserved rather than spent, and conserving water is what lets the cooling continue once the rain stops. The glaze and the moss are, in this sense, two versions of the same move.

Beyond conserving water, the moss carries the contributions the ceramic cannot: it is the living layer that supports urban biodiversity and captures fine dust, and it did so without the night-time heat-trapping that dry moss produces on bare concrete, because the stored water kept it from drying into a dormant, insulating state. So each part of the system answers a different need. The ceramic body cools, the glaze and moss extend how long that cooling lasts, and the moss adds the ecological value that motivated a living facade in the first place.

The clear limit is water. Under the harshest condition tested, a constant 30 °C with no night relief, the saturated porous body gave up most of its store in about a day, and although the outdoor cycle showed real performance to be considerably more forgiving, a single fill is not enough to bridge a multi-day drought with confidence. The most consequential finding of the phase is therefore also its clearest design direction: the system works, but to be genuinely self-sustaining it needs a larger water store or, better, passive rainwater buffering that keeps the store topped up rather than draining it once. The cooling, the longevity mechanisms and the ecological layer are all in place; what remains is to feed them.

06

Closing



Fig. 6.1 — AI-assisted render showing a possible future facade integration



Fig. 6.2 — The prototype with its living layer: moss transplanted into the printed channels of the 3D-printed ceramic facade.



Fig. 6.3 — The same wall low and from the side, the transplanted moss reading as a soft green grain across the printed relief.

A wall printed, fired, and planted by hand. For now the moss is young, pressed into the channels and holding on rather than grown in, but the conditions are set for it to spread. Imagined across a facade, and then across a street, these tiles turn the blank vertical faces of the city into something that cools, holds water, and makes room for life.

Fig. 6.4 — Close-up: moss pressed into the porous valleys where water is held, while the raised peaks stay clear.





Fig. 6.5 — The prototype in the open air: the young transplanted moss tracing the channels of the printed ceramic.



Fig. 6.6 — Visualisation of the system at building scale: the bioreceptive tiles greening a street-level facade (impression, not a built installation).

Up close the wall is honest about what it is: printed ceramic, fired hard, its channels planted by hand. The moss is young and still finding its hold, yet the surface already does the quiet work it was built for, catching water and keeping it where life can reach it.

Set into a street, the same tiles stop reading as a prototype and become a frontage: a wall that holds a little weather, cools the air beside it, and turns a blank brick face into habitat. This is the direction the work points toward, one tile proven, a facade imagined.

Conclusion

6.1

This thesis set out to find how 3D-printed porous ceramics could be used to make modular facade tiles that keep living moss active as a cooling surface and a foothold for urban biodiversity. The answer it arrives at is a material and a method. A ceramic body can be engineered to hold and move water through its own structure, a printed surface can be shaped to catch and slow that water, and the two together can be assembled into a module that feeds the moss growing on it. The work shows that this is achievable, and it maps out clearly the conditions a full working facade would still have to meet.

The argument runs through three scales, and water connects them. At the micro-scale, a white stoneware body with spent coffee grounds and corn starch as pore-formers, fired at 1100 °C, was given the porosity to absorb and store water inside the material, with a verified neutral chemistry and a pore structure designed to resist frost. At the meso-scale, five printed geometries were compared, and a fractal-branching surface showed the most promise, its irregular branching channels holding water at the face and slowing its loss better than the more ordered forms. At the macro-scale, the validated tile was integrated into a module that distributes water across a wall and concentrates it where the moss establishes. Absorb, retain, distribute: each phase handles one part of the same problem, which is keeping water present at the surface for as long as possible. That focus on water follows from the central finding of the project. Moss only cools while it is evaporating, and once it dries it falls dormant and can leave a wall warmer than bare concrete. Sustained moisture is the condition the entire thermal benefit rests on, and the design exists to hold the moss above that threshold through dry periods. The same logic surfaced as a trade-off during testing. A moss-covered, partly glazed surface ran slightly warmer in the moment, yet it held far more water and stayed active far longer, which is the behaviour that matters across a real summer.

What the thesis establishes with confidence is the material. The recipe reaches the bioreceptive porosity window and tests as chemically neutral, and transplanted moss took to the surface and has held for one month and seven days at the time of writing, still healthy and counting. Frost resistance is the one material property not yet demonstrated: it is predicted by the Maage Factor on the assumption of a bimodal pore structure, but that pore distribution was estimated rather than measured and no freeze-thaw test was run, so it remains a design prediction still to be verified. The meso and macro results are first indications. With a single sample per geometry and a cooling study that was specified but not completed in time, the comparison between geometries and the performance of the assembled module are directional, and the facade itself is developed as a conceptual application of a validated tile. The honest position is that the building blocks are proven and the system around them is shown to be plausible and ready to test.

The clearest next steps follow from those gaps. The specified wetness-and-cooling study should be run to measure whether a continuously supplied substrate keeps the moss cooling through a heatwave. A larger basin or a passive rainwater buffer should be added and tested, since the system survives but only keeps cooling for long when it is fed. Geometries should be replicated so the comparison becomes statistical, the biology should be monitored across a full season, and the firing energy of the ceramic should be addressed through measures such as a single combined firing, locally sourced clay, and distributed production. The parametric pipeline gives a way to carry the design forward at any scale once these questions are answered.

Taken as a whole, the project delivers a working prototype, a reproducible method, and a candid account of what a self-sustaining bioreceptive facade demands. A single wall will not undo urban heat, and this one is not finished. What it offers is a credible, tested route toward facades that cool a little, host life, and ask very little to keep doing both.

Discussion

6.2

that is the most useful thing to reflect on. The micro-scale work is the most solid: the ceramic recipe reaches the bioreceptive porosity window while staying pH-neutral (frost resistance is designed in but not yet tested), and it was tested enough to trust. The meso and macro scales are more provisional. Working with single samples per geometry and a validation built around one prototype means the results are indicative rather than conclusive, and they are best read as design direction rather than proof. Spreading the work across all three scales gave the project its range and its narrative, but it came at the cost of depth, and taking one scale further with replicated samples would have produced firmer evidence. The Research-through-Design approach fits this honestly, since its value here was in generating direction and surfacing the right questions rather than settling them.

Several material findings need further work before the recipe can be called finished. The combined mix produced a higher porosity than the individual additives predicted, which suggests the ingredients interact during firing rather than contributing independently. This interaction needs to be investigated and the recipe re-tuned around it. The intended bimodal, dual-porosity structure was inferred rather than measured, and should be verified directly, for example with mercury intrusion porosimetry. The mechanical strength of the fired tiles was also not characterised, and needs testing before any structural claim can be made. Fabrication raised its own practical questions. Printing and finishing the tiles was labour-intensive, and a deployable version would need a less manual, more repeatable production route. Warping during drying and firing meant the tiles did not always sit together perfectly, so dimensional control would have to improve for modules to tile cleanly across a facade.

The environmental cost of the material deserves honest attention. Every tile is fired, and firing is energy-intensive and carries real carbon emissions, which sits awkwardly with a product meant to be ecological. A credible version would have to take this seriously, most directly by firing once rather than twice and combining the bisque and glaze into a single firing, alongside locally sourced clay and distributed production to cut transport.

At the system scale, the central open question is water. The chamber test showed that a single reservoir fill is undersized for a multi-day dry period, so the basin would need to be sized properly for survival, or the module would need an integrated water supply, whether a larger passive store or an active top-up. The cladding and mounting system was only taken to proof-of-concept and would need real development: load transfer, sealing and water hand-off between modules, drainage for frost, and even distribution down a column.

Finally, the project is worth questioning at its premise. The cooling was shown to come from the wet ceramic rather than the moss, so purely from a cooling standpoint it is fair to ask whether moss is the right organism, or whether the same ceramic-and-water system could be adapted into a passive green wall with a water buffer and vascular planting. Moss still earns its place for biodiversity, fine-dust capture and water conservation, but that trade-off should be examined rather than assumed.

Taken together, these points map what a next iteration would need to resolve. The project delivers a validated material and a reproducible method, and just as importantly an honest account of the questions that remain before it could become a deployable facade.

References

06 — REFERENCES

AAWS. (n.d.). Rainwater wall. https://www.aaws.nl/en/rainwaterwall/

Abdallah, Y. K., & Estévez, A. T. (2023). 3D-printed bio-receptive tiles of reaction–diffusion (Gierer-Meinhardt model) for multi-scale algal strains' passive immobilization. *Buildings*, 13(8), 1972.

Abdul Rahman, N. F., Abd. Rahman, H., Wan Azhan, W. M. S., & Yatim, N. H. (2021). Linear shrinkage, strength and porosity of alumina-based ceramic foam with corn starch as pore former. *Solid State Phenomena*, 317, 75–80.

Ahmad, R., Ha, J.-H., & Song, I.-H. (2014). Processing methods for the preparation of porous ceramics. *Journal of Korean Powder Metallurgy Institute*, 21(5), 389–398. https://doi.org/10.4150/KPMI.2014.21.5.389

Alfa, J. I., & Adedoyin, S. O. (2024). A review of urban heat islands: Causes, effects, and mitigation. *International Journal of Agricultural Economics, Management and Development*, 12(1), 83–95.

Archello. (n.d.). LivePanel living wall system with interchangeable plant cassettes. https://archello.com/product/livepanel-living-wall-system-with-interchangeable-plant-cassettes

Asian Development Bank. (2024). How cities can combat extreme heat using nature-based solutions. *Development Asia*.

Asian, A. S., & Mutlu Avinç. (2022). A biomimetic approach to water harvesting strategies. *International Journal of Built Environment and Sustainability*, 9(3), 47–60

Beckett, R. (2021). Probiotic design. *The Journal of Architecture*, 26(1), 6–31.

Bio-receptive architectures: The synthesis of macro-geometry, micro-climatic engineering, and algorithmic design for urban biodiversity. (n.d.).

Bio-receptive concrete – Critical playground. (2024).

Bone, J. R., Stafford, R., Hall, A. E., & Herbert, R. J. H. (2022). The intrinsic primary bio-receptivity of concrete in the coastal environment – A review. *Developments in the Built Environment*, 10, 100078.

Bose, S., Akdogan, E. K., Balla, V. K., Ciliveri, S., Colombo, P., Franchin, G., Ku, N., Kushram, P., Niu, F., Pelz, J., Rosenberger, A., Safari, A., Seeley, Z., Trice, R. W., Vargas-Gonzalez, L., Youngblood, J. P., & Bandyopadhyay, A. (2024). 3D printing of ceramics: Advantages, challenges, applications, and perspectives. *Journal of the American Ceramic Society*, 107(5). https://doi.org/10.1111/jace.20043

Bouzerara, F., Harabi, A., & Boudaira, B. (2009). Preparation and characterization of macroporous ceramic supports for membranes. *Physics Procedia*, 2, 1449–1453.

Comprehensive assessment of the impact of green roofs and walls on building energy performance. (2024). *Energies*, 17, 5160.

Crawford, A., Diniz, N., Beckett, R., Vanucchi, J., & Swackhamer, M. (Eds.). (2023). *Habits of the Anthropocene: Scarcity and abundance in a post-material economy*. Proceedings of the 43rd Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA).

Crawford, A., In-na, P., Caldwell, G., Armstrong, R., & Bridgens, B. (2022). Clay 3D printing as a bio-design research tool: Development of photosynthetic living building components. *Architectural Science Review*, 65(3), 185–195.

Cruz, M., & Beckett, R. (2016). Bio-receptive design: A novel approach to biodigital materiality. *Architectural Research Quarterly*, 20(1), 51–64. https://doi.org/10.1017/S1359135516000130

Cultrone, G., Sebastián, E., Elert, K., de la Torre, M. J., Cazalla, O., & Rodríguez-Navarro, C. (2004). Influence of mineralogy and firing temperature on the porosity of bricks. *Journal of the European Ceramic Society*, 24, 547–564.

Dardir, M., Wilson, J., & Berardi, U. (2026). Natural infrastructure for environmental and community resilience: A data-driven decision-making framework for urban development. *ACSA 114th Annual Meeting: Convergence/ Divergence*.

Dunea. (2023). Uitvinding Rainmate volop in ontwikkeling. https://www.dunea.nl/algemeen/nieuws/2023/uitvinding-rainmate-volop-in-ontwikkeling

Flynn, A. J., & Stachurski, Z. H. (2006). Microstructure and properties of stoneware clay bodies. *Clay Minerals*, 41(3), 775–789.

Gorbushina, A. A. (2007). Life on the rocks. *Environmental Microbiology*, 9(7), 1613–1631. https://doi.org/10.1111/j.1462-2920.2007.01301.x

The Green Village. (2023). Weverling passieve groene gevel. https://www.thegreenvillage.org/project/weverling-groenproviders/

The Green Village. (2025). Leafy gevelpaneel. https://www.thegreenvillage.org/project/leafy-gevelpaneel/

Guillitte, O. (1995). Bio-receptivity: A new concept for building ecology studies. *Science of the Total Environment*, 167(1–3), 215–220. https://doi.org/10.1016/0048-9697(95)04582-1

Gunawardena, K., & Steemers, K. (2020). Urban living walls: Reporting on maintenance challenges from a review of European installations. *Architectural Science Review*, 63(6), 526–535. https://doi.org/10.1080/00038628.2020.1738209

Gunawardena, K., & Steemers, K. (n.d.). Living walls in indoor environments. Department of Architecture, University of Cambridge.

Hadi, E. M., & Al-Hussien, H. J. (2017). Manufacturing porous ceramic from Iraqi kaolin by using paper pulp. *Engineering and Technology Journal*, 35, 227–236.

Hammel, E. C., Ighodaro, O. L., & Okoli, O. I. (2014). Processing and properties of advanced porous ceramics: An application based review. *Ceramics International*, 40(10), 15351–15370.

Hartanto, S., & Janasthi, N. (2023). The use of coffee grounds as an additional material in ceramic planter products. *International Seminar on Language, Education, and Culture (ISOLEC)*, 282–286.

Ing, W., & Jessel, E. (n.d.). The green building boom / Why some green walls fail. *The Architects' Journal*.

Jakubovskis, R. (2025). Biophilic façades: The potentiality of bio-receptive concrete. *Buildings*, 15, 3646. https://doi.org/10.3390/buildings1520364

Japan for Sustainability. (2008). Moss greening to curb urban heat island effect.

Jiang, F., Zhang, L., Jiang, Z., Li, C., Cang, D., Liu, X., Xuan, Y., & Ding, Y. (2019). Diatomite-based porous ceramics with high apparent porosity: Pore structure modification using calcium carbonate. *Ceramics International*, 45(5), 6085–6092.

Jordan, M. M., Montero, M. A., Meseguer, S., & Sanfeliu, T. (2008). Influence of firing temperature and mineralogical composition on bending strength and porosity of ceramic tile bodies. *Applied Clay Science*, 42, 266–271.

Kinoshita, H., Yasuda, Y., Nakazono, T., Yasui, K., Yuji, T., Okamura, Y., Kobayashi, T., & Fukuyama, H. (2014). Greening material consisting of a porous ceramic made from waste glass fiber reinforced plastic coated with moss. *International Journal of Innovations in Engineering and Technology (IJJET)*.

Li, S., Wang, C.-A., & Zhou, J. (2013). Effect of starch addition on microstructure and properties of highly porous alumina ceramics. *Ceramics International*, 39, 8833–8839.

Lim, A. C. S., & Lharchi, A. (2023). Parameters for bio-receptivity in 3D printing. In B. Faircloth et al. (Eds.), *Design for climate adaptation* (pp. 709–723). Springer. https://doi.org/10.1007/978-3-031-36320-7_44

Liu, G. (2011). Fabrication of porous ceramics and composites by a novel freeze casting technique [Doctoral dissertation, The University of Birmingham].

07

08

Maciel, F. S., Areias, I. O. R., & Holanda, J. N. F. (2023). Valorization potential of coffee grounds waste as a renewable pore-forming agent to produce low-cost porous ceramic support. *Research, Society and Development*, 12(2), e22612240023.

Manso, M., & Castro-Gomes, J. (2015). Green wall systems: A review of their characteristics. *Renewable and Sustainable Energy Reviews*, 41, 863–871. https://doi.org/10.1016/j.rser.2014.07.203

Matser, D. (2023). Enhancing the bio-receptivity of ceramic RainReefs. Utrecht University & Urban Reef.

Miller, A. Z., Rogerio-Candeleria, M. A., Dionísio, A., Macedo, M. F., & Saiz-Jiménez, C. (2012). Evaluación de la influencia de la rugosidad superficial sobre la colonización epilítica de calizas mediante técnicas sin contacto. *Materiales de Construcción*, 62(307), 411-424.

Mustafa, K. F., Prieto, A., & Ottelé, M. (2021). The role of geometry on a self-sustaining bio-receptive concrete panel for facade application. *Sustainability*, 13(13), 7453.

Nigay, P. M., Sani, R., Cutard, T., & Nzihou, A. (2017). The impact of heat treatment on the microstructure of a clay ceramic and its thermal and mechanical properties. *Materials Science and Engineering: A*, 708, 375–386.

Nuruzzaman, M. (2015). Urban heat island: Causes, effects and mitigation measures — A review. *International Journal of Environmental Monitoring and Analysis*, 3(2), 67–73.

Ochen, W., Dujanga, F. M., & Oruru, B. (2019). Effect of quartz particle size on sintering behavior and flexural strength of porcelain tiles made from raw materials in Uganda. *JOJ Material Science*, 5(3), 555661.

Pabst, W., Sandoval, M. L., Talou, M. H., & Tomba Martinez, A. G. (2017). Porous cordierite-based ceramics processed by starch consolidation casting – Microstructure and high-temperature mechanical behavior. *Ceramics International*.

Parveen, S., & Tariq, M. (2023). Bio-receptive building materials for urban ecology. *New Environmentally-Friendly Materials*, 2(1), 41–55.

Pathak, K., & Brahmandam, A. K. S. V. (2025). Innovative construction technologies for green and climate-resilient housing: The potential of moss concrete.

Pentecost, A., & Whitton, B. A. (2012). Subaerial cyanobacteria. In B. Whitton (Ed.), *Ecology of Cyanobacteria II: Their diversity in space and time* (pp. 291–316). Springer. https://doi.org/10.1007/978-94-007-3855-3_10

Pollini, B., Contardo, T., & Paciotti, D. (n.d.). Generative and bio-receptive design. UNICAM.

Pollini, B., & Rognoli, V. (2021). Enhancing living/non-living relationships through designed materials. Responsible biotechnologies and biodesign for the built environment. *International Conference Construction, Energy, Environment & Sustainability CEES 2021*. Coimbra.

Proctor, M. C., Oliver, M. J., Wood, A. J., Alpert, P., Stark, L. R., Cleavitt, N. L., & Mishler, B. D. (2007). Desiccation-tolerance in bryophytes: A review. *The Bryologist*, 110(4), 595–621.

Radić, M., Brković Dodig, M., & Auer, T. (2019). Green facades and living walls — A review establishing the classification of construction types and mapping the benefits. *Sustainability*, 11, 4579.

Raghani, S. (2025). Moss Concrete: An innovative material for sustainable architecture. *Tuijin Jishu / Journal of Propulsion Technology*, 46(1).

Rotondi, C., Gironi, C., Ciuffo, D., Diana, M., & Lucibello, S. (2024).

Salleh, S. Z., Awang Kechik, A., Yusoff, A. H., Taib, M. A. A., Mohamad Nor, M., Mohamad, M., & Mohamed, J. J. (2021). Recycling food, agricultural, and industrial wastes as pore-forming agents for sustainable porous ceramic production: A review. *Journal of Cleaner Production*, 306, 127264.

Secrid. (2023). Secrid talent podium: Respyre. https://secrid.com/en-us/stories/secrid-talent-podium-respyre/

Severi, A. (2023). 3D printing "living" bio-based materials with Urban Reef. WASP.

Silva, K. R., Menezes, R. R., Campos, L. F. A., & Santana, L. N. L. (2022). A review on the production of porous ceramics using organic and inorganic industrial waste. *Cerâmica*, 68(387), 270–284.

Sochůrková, P., Vaško, I., Tsikolia, S., Sviták, D., Latour, M., & Oskam, P. (2023). Bio-receptivity as a factor of additive fabrication: Material research on 3D printed structures affecting urban spaces. *eCAADe 41*, 2, 115–124. https://doi.org/10.52842/conf.ecaade.2023.2.115

Soza Ruiz, P., Navarro Galleguillos, L., & García Alvarado, R. (2025). On the bio-recepte properties of 3D printed clay tiles: A methodological model for enhancing moss colonization through geometric design. *SiGraDi 2025*.

Stohl, L., Manninger, T., von Werder, J., Dehn, F., Gorbushina, A., & Meng, B. (2023). Bio-receptivity of concrete: A review. *Journal of Building Engineering*, 76, 107201.

Stohl, L., Manninger, T., Dehn, F., & von Werder, J. (2025). Understanding bio-receptivity of concrete: Material design and characterization. *Materials and Structures*, 58(335).

Stryszewska, T., & Kanka, S. (2022). Evaluation of frost impact on traditional ceramic building materials utilized in facing walls. *Materials*, 15(17), 5858. *Sustainability Directory*. (2026). Evapotranspiration cooling. Tang, H., Liang, X., Cao, T., Salley, S. O., McAllister, J. P., & Ng, K. Y. S. (2006). Influence of surface roughness and hydrophobicity on bacterial adhesion and colonization. 2006 AIChE Annual Meeting.

The science of bio-receptive substrates: Topographical metrology, rugosity index, and the mechanics of biological colonization. (n.d.).

Thermophysical dynamics of urban microclimates: Evapotranspirative cooling, green infrastructure systems, and resource constraints. (n.d.).

Tietz, H. (2021). Foreground for mosses: Designing 3D printed clay bryobricks to enhance the built environment. *ASLA 2021 Student Awards*.

Tomba Martinez, A. G., Sandoval, M. L., & Camerucci, M. A. (2012). High-temperature mechanical behavior of cordierite-based porous ceramics prepared by modified cassava starch thermogelation.

TU Delft. (2023). Respyre blaast met groen beton de stad nieuw leven in. https://www.tudelft.nl/innovatie-impact/pioneering-tech/articles/respyre-blaast-met-groen-beton-de-stad-nieuw-leven-in

Veeger, M. I. A., Ottelé, M., & Prieto Hoces, A. I. (2021). Making bio-receptive concrete: Formulation and testing of bio-receptive concrete mixtures. *Journal of Building Engineering*, 44, 102545. https://doi.org/10.1016/j.jobe.2021.102545

Veeger, M. I. A., Ottelé, M., & Jonkers, H. M. (2026a). How moss affects urban temperatures: The effects of moss on the thermal dynamics of an urban cementitious surface. *Science of the Total Environment*, 1034, 181830. https://doi.org/10.1016/j.scitotenv.2026.181830

Veeger, M., Veenendaal, E. M., Limpens, J., Ottelé, M., & Jonkers, H. M. (2025). Moss species for bio-receptive concrete: A survey of epilithic urban moss communities and their dynamics. *Ecological Engineering*, 212, 107502.

Wilson, K. (2021, September 20). Biden's 'Heat Island' Strategy Ignores Cars. *Streetsblog USA*. https://usa.streetsblog.org/2021/09/20/bidens-heat-island-strategy-ignores-cars

Zanelli, C., Raimondo, M., Dondi, M., Guarini, G., & Tenorio Cavalcante, P. M. (2004). Sintering mechanisms of porcelain stoneware tiles. *Qualicer 2004*, P.GI-248.

Zheng, Z., Fong, C. S., Aghamohammadi, N., & Law, Y. K. (2026). A systematic review of urban heat island (UHI) impacts and mitigation: Health, equity, and policy. *Systems*, 14(1), 82. https://doi.org/10.3390/systems14010082

Sample catalogue

Apparent porosity and water absorption per fired sample (ASTM C20-00 and ISO 10545-3). Values as recorded; decimal comma.

Sample	SCG	Diat	Starch	Grog	Firing	m before	m after	Loss	Susp	Vol	m ASTM	App. por.	Water abs	Por. ISO	m ISO	Note
Control — no additive																
2-B0-A	0%	0%	0%	25%	1000		17,24		10,49	8,91	19,4	24,24%	12,53%			
2-B0-B	0%	0%	0%	25%	1000		20,83		12,55	11,05	23,6	25,07%	13,30%			
2-B0-C	0%	0%	0%	25%	900		15,75		9,5	8,5	18	26,47%	14,29%			
2-B0-D	0%	0%	0%	25%	900		16,36		9,99	9,01	19	29,30%	16,14%			
B0-A	0%	0%	0%	25%	900	488,47	460,52	5,72%	260	273	533,0	26,55%	15,74%	12,26%	517	
B0-A1	0%	0%	0%	25%	1100		25,25		14,95	12,45	27,4	17,27%	8,51%			
B0-B	0%	0%	0%	25%	1000	486,32	458,4	5,74%	255	275	530	26,04%	15,62%			
B0-B2	0%	0%	0%	25%	1100		19,11		11,3	9,3	20,6	16,02%	7,80%			
B0-C1	0%	0%	0%	25%	1100	472,18	213	54,89%	104	126	230	13,49%	7,98%			
B0-C2	0%	0%	0%	25%	1100	472,18	85	82,00%	42	52	94	17,31%	10,59%			
B0-C3	0%	0%	0%	25%	1100		26,06		15,33	13,07	28,4	17,90%	8,98%			
B0-D	0%	0%	0%	25%	1100		22,1		13,08	10,92	24	17,40%	8,60%			
B0-E	0%	0%	0%	25%	1100		19,68		11,65	9,75	21,4	17,64%	8,74%			
B0-G	0%	0%	0%	25%	1100		19,66		11,62	9,78	21,4	17,79%	8,85%			
B0-H	0%	0%	0%	25%	1100		22,98		13,49	11,11	24,6	14,58%	7,05%			
Spent coffee grounds (SCG)																
2-C10-A	10%	0%	0%	25%	1100		22,63		13,09	14,71	27,8	35,15%	22,85%			
2-C10-B	10%	0%	0%	25%	1100		22,26		12,85	14,35	27,2	34,43%	22,19%			
2-C10-C	10%	0%	0%	25%	1000		21,79		12,59	15,61	28,2	41,06%	29,42%			
2-C10-D	10%	0%	0%	25%	900		18,53		10,53	13,87	24,4	42,32%	31,68%			
2-C15-A	15%	0%	0%	25%	1100		20,22		12,39	14,61	27	46,41%	33,53%			
2-C15-B	15%	0%	0%	25%	1000		17,18		10,52	13,68	24,2	51,32%	40,86%			
2-C15-C	15%	0%	0%	25%	1000		21,96		13,45	17,75	31,2	52,06%	42,08%			
2-C15-D	15%	0%	0%	25%	900		22,07		13,22	18,18	31,4	51,32%	42,27%			
2-C5-A	5%	0%	0%	25%	1100		14,61		8,21	8,39	16,6	23,72%	13,62%			
2-C5-B	5%	0%	0%	25%	1000		11,05		6,46	6,94	13,4	33,86%	21,27%			
2-C5-C	5%	0%	0%	25%	1100		9,2		5,28	5,32	10,6	26,32%	15,22%			
2-C5-D	5%	0%	0%	25%	1000		8,55		5,04	5,36	10,4	34,51%	21,64%			
2-C5-E	5%	0%	0%	25%	900		10,83		6,22	6,98	13,2	33,95%	21,88%			
2-C5-F	5%	0%	0%	25%	900		10,9		6,17	7,03	13,2	32,72%	21,10%			
C10-A	10%	0%	0%	25%	900	359,22	315,73	12,11%	233	172,2	405,2	51,96%	28,34%	20,67%	381	
C10-B	10%	0%	0%	25%	900	370,37	324,78	12,31%	226	192,4	418,4	48,66%	28,83%	20,08%	390	
C10-C	10%	0%	0%	25%	1000	355,24	311,45	12,33%	216	172	388	44,51%	24,58%			
C10-D	10%	0%	0%	25%	1000	363,61	318,75	12,34%	221	175,5	396,5	44,30%	24,39%			
C10-E	10%	0%	0%	25%	1100	367,27	314,61	14,34%	203	174,1	377,1	35,89%	19,86%			
C15-A	15%	0%	0%	25%	900	198,16	122	38,43%	101	75	176	72,00%	44,26%	86,89%	228	Broken during transport, unfired
C15-B	15%	0%	0%	25%	900	379,00	312,86	17,45%	258	190,6	448,6	71,22%	43,39%	36,48%	427	
C15-C	15%	0%	0%	25%	1000	383,83	316,98	17,42%	249	193,8	442,8	64,92%	39,69%			
C15-D	15%	0%	0%	25%	1100	370,55	305,58	17,53%	213	181	394	48,85%	28,94%			
C15-E	15%	0%	0%	25%	1100	348,98	288,42	17,35%	207	172,2	379,2	52,72%	31,47%			
C5-A	5%	0%	0%	25%	900	320,40	292,13	8,82%	181	160	341	30,54%	16,73%	17,76%	344	Broken during transport, unfired
C5-B	5%	0%	0%	25%	900	388,46	353,73	8,94%	230	204,2	434,2	39,41%	22,75%	15,34%	408	Damage due to airbubbles while printing
C5-C	5%	0%	0%	25%	1000	386,54	352,07	8,92%	223	198	421	34,81%	19,58%			
C5-D	5%	0%	0%	25%	1100	388,33	353,35	9,01%	194	195,4	389,4	18,45%	10,20%			
Diatomite																
2-D10-A	0%	10%	0%	25%	1100		22,42		13,5	11,1	24,6	19,64%	9,72%			
2-D10-B	0%	10%	0%	25%	900		22,91		13,71	13,09	26,8	29,72%	16,98%			
2-D15-A	0%	15%	0%	25%	900		24,21		14,42	14,38	28,8	31,92%	18,96%			
2-D5-A	0%	5%	0%	25%	1100		16,79		10,11	8,09	18,2	17,43%	8,40%			
2-D5-B	0%	5%	0%	25%	900		15,93		9,66	8,94	18,6	29,87%	16,76%			
2-D5-C	0%	5%	0%	25%	900		14,94		9,08	8,32	17,4	29,57%	16,47%			
D10-A	0%	10%	0%	25%	1000 1s	41,59	39	6,23%	22,2	22,6	44,8	25,66%	14,87%			
D10-B	0%	10%	0%	25%	900 1s	39,09	36,71	6,09%	20,2	22,8	43	27,59%	17,13%			
D10-C	0%	10%	0%	25%	900 1s	40,99	38,48	6,12%	22,1	23,9	46	31,46%	19,54%			
D10-D	0%	10%	0%	25%	1000	31,42	29,44	6,30%	16,8	17,2	34	26,51%	15,49%			
D10-E	0%	10%	0%	25%	1000	31,82	29,82	6,29%	17,4	17	34,4	26,94%	15,36%			
D10-F	0%	10%	0%	25%	1000	32,7	30,66	6,24%	18,52	16,58	35,1	26,78%	14,48%			
D10-H	0%	10%	0%	25%	1100	34,28	32,08	6,42%	15,6	18,6	34,2	11,40%	6,61%			
D10-I	0%	10%	0%	25%	1100	28,13	26,33	6,40%	12,7	15,5	28,2	12,06%	7,10%			
D15-A	0%	15%	0%	25%	1000 1s	29,9	28	6,35%	15,9	16,6	32,5	27,11%	16,07%			
D15-B	0%	15%	0%	25%	900 1s	30,77	28,86	6,21%	17	18	35	34,11%	21,28%			
D15-C	0%	15%	0%	25%	900 1s	34,73	32,55	6,28%	18,9	20,1	39	32,09%	19,82%			
D15-D	0%	15%	0%	25%	1000	31,85	29,75	6,59%	17,23	16,77	34	25,34%	14,29%			
D15-E	0%	15%	0%	25%	1000	28,16	26,32	6,53%	14,3	15,9	30,2	24,40%	14,74%			
D15-F	0%	15%	0%	25%	1000	29,48	27,56	6,51%	15,2	16,3	31,5	24,17%	14,30%			
D15-G	0%	15%	0%	25%	1100	26,55	24,78	6,67%	12,8	13,7	26,5	12,55%	6,94%			
D15-H	0%	15%	0%	25%	1100	31,47	29,37	6,67%	14,6	16,6	31,2	11,02%	6,23%			
D15-I	0%	15%	0%	25%	1100	32,91	30,73	6,62%	15,1	17,6	32,7	11,19%	6,41%			
D5-A	0%	5%	0%	25%	1000 1s	39,52	37,2	5,87%	20,9	21,8	42,7	25,23%	14,78%			
D5-B	0%	5%	0%	25%	900 1s	37,24	35,06	5,85%	19,9	21,1	41	28,15%	16,94%			
D5-C	0%	5%	0%	25%	1000	35,7	33,58	5,94%	20,26	18,24	38,5	26,97%	14,65%			
D5-D	0%	5%	0%	25%	1000	32,96	30,96	6,07%	17,3	18,1	35,4	24,53%	14,34%			
D5-E	0%	5%	0%	25%	1000	38,45	36,16	5,96%	20,2	21	41,2	24,00%	13,94%			
D5-F	0%	5%	0%	25%	1100	35,7	33,54	6,05%	17,3	18,5	35,8	12,22%	6,74%			
D5-H	0%	5%	0%	25%	1100	32,81	30,8	6,13%	14,7	18,2	32,9	11,54%	6,82%			
Corn starch																
S10-A	0%	0%	10%	25%	1100	12,92	11,07	14,32%	6,54	6,46	13	29,88%	17,43%			
S10-AB	0%	0%	10%	25%	900		13,94		7,7	8,9	16,6	29,89%	19,08%			

Sample catalogue

Apparent porosity and water absorption per fired sample (ASTM C20-00 and ISO 10545-3). Values as recorded; decimal comma.

Sample	SCG	Diat	Starch	Grog	Firing	m before	m after	Loss	Susp	Vol	m ASTM	App. por.	Water abs	Por. ISO	m ISO	Note
S10-AC	0%	0%	10%	25%	900		15,67		8,68	10,12	18,8	30,93%	19,97%			
S10-AD	0%	0%	10%	25%	900		15,53		8,58	9,82	18,4	29,23%	18,48%			
S10-AE	0%	0%	10%	25%	900		12,12		6,78	7,62	14,4	29,92%	18,81%			
S10-B	0%	0%	10%	25%	1100	16,99	14,61	14,01%	8,46	8,54	17	27,99%	16,36%			
S10-C	0%	0%	10%	25%	1100	20,61	17,69	14,17%	10,24	10,36	20,6	28,09%	16,45%			
S10-D	0%	0%	10%	25%	1100	21,08	18,1	14,14%	10,43	10,87	21,3	29,44%	17,68%			
S10-E	0%	0%	10%	25%	1100	14,56	12,49	14,22%	7,29	7,31	14,6	28,86%	16,89%			
S10-F	0%	0%	10%	25%	1000		16,4		9,16	10,24	19,4	29,30%	18,29%			
S10-G	0%	0%	10%	25%	1000		13,5		7,66	8,34	16	29,98%	18,52%			
S10-H	0%	0%	10%	25%	1000		14,1		7,96	8,84	16,8	30,54%	19,15%			
S10-I	0%	0%	10%	25%	1000		14,7		8,33	9,27	17,6	31,28%	19,73%			
S13-A	0%	0%	13%	25%	1100		18,72		10,86	10,74	21,6	26,82%	15,38%			
S13-B	0%	0%	13%	25%	1100		18,85		11,06	10,94	22	28,79%	16,71%			
S13-C	0%	0%	13%	25%	1100		22,93		13,33	13,47	26,8	28,73%	16,88%			
S15-A	0%	0%	15%	25%	1100	17,76	14,63	17,62%	8,54	8,86	17,4	31,26%	18,93%			
S15-AB	0%	0%	15%	25%	900		12,78		7,05	8,35	15,4	31,38%	20,50%			
S15-AC	0%	0%	15%	25%	900		13,41		7,29	9,11	16,4	32,82%	22,30%			
S15-AD	0%	0%	15%	25%	900		11,97		6,57	8,23	14,8	34,39%	23,64%			
S15-AE	0%	0%	15%	25%	900		9,35		5,11	6,49	11,6	34,67%	24,06%			
S15-B	0%	0%	15%	25%	1100	16,42	13,53	17,60%	7,95	8,05	16	30,68%	18,26%			
S15-C	0%	0%	15%	25%	1100	16,25	13,42	17,42%	7,81	8,09	15,9	30,66%	18,48%			
S15-D	0%	0%	15%	25%	1100	15	12,35	17,67%	7,24	7,46	14,7	31,50%	19,03%			
S15-E	0%	0%	15%	25%	1100	17,32	14,28	17,55%	8,22	8,78	17	30,98%	19,05%			
S15-F	0%	0%	15%	25%	1000		11,7		6,65	7,75	14,4	34,84%	23,08%			
S15-G	0%	0%	15%	25%	1000		11,2		6,32	7,48	13,8	34,76%	23,21%			
S15-H	0%	0%	15%	25%	1000		10,4		5,94	6,86	12,8	34,99%	23,08%			
S15-I	0%	0%	15%	25%	1000		13,6		7,69	9,21	16,9	35,83%	24,26%			
S15-J	0%	0%	15%	25%	1000		11,9		6,7	7,9	14,6	34,18%	22,69%			
S2-A	0%	0%	2%	25%	1100		23,64		13,81	11,79	25,6	16,62%	8,29%			
S2-B	0%	0%	2%	25%	1100		21,47		12,61	10,79	23,4	17,89%	8,99%			
S2-C	0%	0%	2%	25%	1100		20,5		12	10,4	22,4	18,27%	9,27%			
S5-A	0%	0%	5%	25%	1100	15,19	13,64	10,20%	7,95	7,45	15,4	23,62%	12,90%			
S5-AB	0%	0%	5%	25%	900		12,33		7,11	7,49	14,6	30,31%	18,41%			
S5-AC	0%	0%	5%	25%	900		13,56		7,78	8,02	15,8	27,93%	16,52%			
S5-AE	0%	0%	5%	25%	900		14,73		8,51	8,69	17,2	28,42%	16,77%			
S5-AF	0%	0%	5%	25%	900		11,57		6,67	6,93	13,6	29,29%	17,55%			
S5-B	0%	0%	5%	25%	1100	15,59	14	10,20%	8,15	7,55	15,7	22,52%	12,14%			
S5-C	0%	0%	5%	25%	1100	15,81	14,22	10,06%	8,25	7,75	16	22,97%	12,52%			
S5-D	0%	0%	5%	25%	1100	14,85	13,33	10,24%	7,75	7,25	15	23,03%	12,53%			
S5-E	0%	0%	5%	25%	1100	15,02	13,49	10,19%	7,93	7,57	15,5	26,55%	14,90%			
S5-F	0%	0%	5%	25%	1000		14,1		8,25	8,35	16,6	29,94%	17,73%			
S5-G	0%	0%	5%	25%	1000		11,9		6,94	7,06	14	29,75%	17,65%			
S5-H	0%	0%	5%	25%	1000		13,3		7,77	7,83	15,6	29,37%	17,29%			
S5-I	0%	0%	5%	25%	1000		13,1		7,64	7,76	15,4	29,64%	17,56%			
Grog / chamotte																
G32,5-A	0%	0%	0%	32,5%	1100		26,4		15,61	13,19	28,8	18,20%	9,09%			
G32,5-B	0%	0%	0%	32,5%	1100		24,12		14,19	12,01	26,2	17,32%	8,62%			
G32,5-C	0%	0%	0%	32,5%	1100		21,42		12,73	10,87	23,6	20,06%	10,18%			
G40-A	0%	0%	0%	40,0%	1100		21,97		12,92	11,08	24	18,32%	9,24%			
G40-B	0%	0%	0%	40,0%	1100		20,36		12,12	10,28	22,4	19,84%	10,02%			
G40-C	0%	0%	0%	40,0%	1100		20,05		11,83	10,17	22	19,17%	9,73%			
High-SCG test print																
TP-SCG16	16%	0%	0%	40,0%	1100		921,43		666	536	1202	52,35%	30,45%			
Final mix (SCG + corn starch)																
Mix-A		0%		40%	1100		12,4		7,9	7,7	15,6	41,56%	25,81%			
Mix-B		0%		40%	1100		19,97		11,86	12,94	24,8	37,33%	24,19%			
Mix-C		0%		40%	1100		13,1		7,88	8,52	16,4	38,73%	25,19%			

Geometry and surface-water behaviour

3D-scan roughness, capillary uptake (sorptivity) and surface drying of the printed geometries.

Surface roughness (3D scan)

Geometry	Rugosity true/projected	Sz mm
Tree Bark (G01)	3.6235	22,38
False Diffusion (G02)	4.0412	16,15
Spherical Pattern (G03)	4.5657	33,18
Woven Fabric (G04)	4.1516	18,43
Fractal Branching (G05)	4.3123	19,10

Capillary uptake (sorptivity)

Cumulative specimen mass (g) absorbed from the base over time; includes two thickness variants of the diffusion pattern.

Geometry	Area cm ²	base g	1m	5m	10m	30m	60m	2h	3h45	21h15	Vol cm ³
Tree Bark	48,09	937,33	953,53	959,75	963,67	971,76	979,74	992,01	1006,82	1072,06	649,2
False Diffusion	49,40	942,28	974,25	985,27	991,46	1005,30	1018,36	1036,02	1057,22	1100,99	666,9
Diffusion thick	54,8	1045,57	1071,62	1080,62	1085,67	1093,87	1102,95	1113,43	1127,80	1189,52	739,8
Diffusion thin	51,00	841,97	866,73	874,57	879,29	888,38	897,89	909,07	924,62	972,28	688,5
Spherical Pattern	47,92	884,40	911,70	924,23	930,08	943,02	955,25	970,13	990,37	1023,50	646,9
Woven Fabric	38,40	892,17	917,20	927,89	932,05	940,81	948,16	957,91	969,81	1017,44	518,4
Fractal Branching	44,87	1017,34	1052,64	1069,92	1077,35	1092,67	1106,73	1123,55	1145,67	1188,38	605,7

Surface drying of the five geometries

Specimen mass over time (g)

Geometry	Dry mass	Sat.	0 10:00	0.25 10:15	0.5 10:30	1 11:00	2 12:00	3 13:00	4 14:00	5 15:00	6 16:00	7 17:00	8.25 18:00
Woven Fabric	38,40	1013,53	1014,45	1011,52	1009,70	1003,25	996,83	985,97	979,86	973,73	965,65	959,86	953,76
Tree Bark	48,09	1065,41	1066,81	1064,58	1062,52	1058,85	1051,59	1045,09	1038,87	1026,21	1020,54	1010,52	1000,52
False Diffusion	49,40	1092,18	1094,34	1092,12	1090,18	1080,69	1073,34	1053,98	1046,08	1029,83	1021,98	1007,49	999,20
Spherical Pattern	47,92	1026,72	1028,10	1022,95	1020,41	1015,35	1006,67	998,27	990,32	982,11	973,98	965,86	955,84
Fractal Branching	44,87	1179,4	1182,01	1179,35	1176,79	1172,07	1161,45	1152,56	1142,70	1130,89	1121,35	1111,54	1099,89

Relative water loss (% of total absorbed)

Geometry	0	0.25	0.5	1	2	3	4	5	6	7	8.25
Woven Fabric	0.00%	0.30%	0.49%	1.15%	1.81%	2.92%	3.54%	4.17%	5.00%	5.59%	6.22%
Tree Bark	0.00%	0.22%	0.42%	0.78%	1.49%	2.13%	2.74%	3.99%	4.54%	5.53%	6.51%
False Diffusion	0.00%	0.21%	0.40%	1.31%	2.01%	3.86%	4.62%	6.17%	6.92%	8.31%	9.10%
Spherical Pattern	0.00%	0.53%	0.78%	1.30%	2.19%	3.04%	3.85%	4.69%	5.52%	6.35%	7.37%
Fractal Branching	0.00%	0.23%	0.46%	0.87%	1.81%	2.59%	3.46%	4.50%	5.33%	6.20%	7.22%

Module longevity and pH

Moss-module drying at a constant 30 °C, and the aqueous-suspension pH test (ISO 787-9).

Phase 3 — longevity of the moss module (constant 30 °C, 19 h)

Specimen mass over time (g)

Specimen	Dry mass	Sat.	0 0m	0.43 25m	0.57 34m	1.02 1h01	1.62 1h37	2.13 2h08	2.67 2h40	3.47 3h28	4.22 4h13	4.80 4h48	19.07 19h04
Moss Module	1333.59	1933.63	1933.63	1900.63	1895.95	1882.92	1870.15	1857.61	1845.55	1825.9	1804.67	1789.27	1457.85

Relative water loss (% of total absorbed)

Specimen	0	0.43	0.57	1.02	1.62	2.13	2.67	3.47	4.22	4.80	19.07
Moss Module	0.00%	5.50%	6.28%	8.45%	10.58%	12.67%	14.68%	17.95%	21.49%	24.06%	79.29%

Phase 1 — pH (chemical compatibility, ISO 787-9)

All ceramic samples returned pH 6 at every interval and reading, identical to the distilled-water baseline — no measurable leaching.

Interval	Sample	Reading 1	Reading 2	Reading 3
Baseline · 0 min	Distilled water	6	6	6
5 minutes	Ceramic sample 1	6	6	6
	Ceramic sample 2	6	6	6
	Ceramic sample 3	6	6	6
1 hour	Ceramic sample 1	6	6	6
	Ceramic sample 2	6	6	6
	Ceramic sample 3	6	6	6
24 hours	Ceramic sample 1	6	6	6
	Ceramic sample 2	6	6	6
	Ceramic sample 3	6	6	6

Parametric G-code generator

Python script that grows the facade geometry, writes the printer G-code directly, and builds the interactive HTML preview. Listed in full; read top-down, left column then right.

```

1 import math
2 import json
3 from shapely.geometry import MultiLineString, LineString, Polygon, box
4
5 # =====
6 # 1. MATRIX GENERATION SETTINGS
7 # =====
8 MATRIX_ROWS = 6
9 MATRIX_COLS = 6
10
11 # =====
12 # 2. CLAY PARAMETERS
13 # =====
14 NOZZLE_DIA = 6.0
15 LAYER_HEIGHT = 3.2
16 Z_MAX = 75.0
17
18 E_MULTIPLIER = 0.2118
19 FEEDRATE = 2100
20 XY_TRAVEL_FEEDRATE = 6000
21 Z_FEEDRATE = 12000
22
23 # =====
24 # 3. PARAMETRIC SETTINGS
25 # =====
26 TILE_WIDTH = 150.0
27 TOTAL_WIDTH = MATRIX_COLS * TILE_WIDTH
28
29 GLOBAL_SPAN_POINTS = [25 + 50 * i for i in range(MATRIX_COLS * 3)]
30
31 BASE_BRANCH_LENGTH = 10.0
32 BASE_BRANCH_ANGLE = 50.0
33 MAX_DEPTH = 4
34 Y_MAX = 40.0
35
36 X_L = -(NOZZLE_DIA / 2.0)
37 X_R = TILE_WIDTH + (NOZZLE_DIA / 2.0)
38
39 FRONT_WALL_DEPTH = 5.0
40 NOTCH_WIDTH = 12.0
41 BACK_WALL_DEPTH = 10.0
42 NOTCH_INSET = 15.0
43
44 BACK_CORRUGATION_AMP = 8 # mm, half-amplitude of back wall corrugation ridges
45 BACK_CORRUGATION_WAVE = 37.5 # mm, wavelength of corrugation
46 BACK_CORRUGATION_PTS = 10 # sample points per wavelength
47
48 EYELET_INSET = 22.0
49 EYELET_DIAMETER = 20.0
50 EYELET_NECK = 3.0
51 EYELET_Z_HEIGHT = 35.0
52
53 BOTTOM_SOLID_LAYERS = 1
54 NUM_LAYERS = int(Z_MAX / LAYER_HEIGHT)
55
56 Y_WAVE_BASE = 0.0
57 Y_NOTCH_TOP = Y_WAVE_BASE - FRONT_WALL_DEPTH
58 Y_NOTCH_BOTTOM = Y_NOTCH_TOP - NOTCH_WIDTH
59 Y_BACK = Y_NOTCH_BOTTOM - BACK_WALL_DEPTH
60
61 # =====
62 # 4. BRANCHING & GLOBAL SLICING LOGIC
63 # =====
64 def grow_branch(x, y, angle_deg, length, branch_angle, depth, segments):
65     if depth == 0: return
66     angle_rad = math.radians(angle_deg)
67     x2 = x + length * math.cos(angle_rad)
68     y2 = y + length * math.sin(angle_rad)
69
70     if y2 > Y_MAX:
71         ratio = (Y_MAX - y) / (y2 - y); x2 = x + (x2-x)*ratio; y2 = Y_MAX
72     if x2 < 0.0:
73         ratio = (0.0 - x) / (x2-x); x2 = 0.0; y2 = y + (y2-y)*ratio
74     elif x2 > TOTAL_WIDTH:
75         ratio = (TOTAL_WIDTH - x) / (x2-x); x2 = TOTAL_WIDTH; y2 = y + (y2-y)*ratio
76
77     segments.append((x, y), (x2, y2))
78     new_length = length * 0.95
79     grow_branch(x2, y2, angle_deg + branch_angle, new_length, branch_angle, depth-1, segments)
80     grow_branch(x2, y2, angle_deg - branch_angle, new_length, branch_angle, depth-1, segments)
81
82     def distance(p1, p2):
83         return math.sqrt((p2[0]-p1[0])**2 + (p2[1]-p1[1])**2)
84
85     def reorder_path(coords, target_pt):
86         if not coords: return coords
87         best_idx = min(range(len(coords)), key=lambda i: distance(coords[i], target_pt))
88         return coords[best_idx:-1] + coords[:best_idx] + [coords[best_idx]]
89
90     GLOBAL_POLYGONS = {}
91
92     def precompute_global_polygons():
93         print(f"Precomputing continuous organic master geometry across {TOTAL_WIDTH}mm...")
94         total_layers_across_wall = MATRIX_ROWS * NUM_LAYERS
95
96         for global_layer_idx in range(total_layers_across_wall):
97             # Match the inspiration exactly: one sin per layer for angle, one cos for length.
98             angle_swing = math.sin(global_layer_idx * 0.15) * 15.0
99             current_angle = BASE_BRANCH_ANGLE + angle_swing
100             length_pulse = math.cos(global_layer_idx * 0.10) * 3.0
101             current_length = BASE_BRANCH_LENGTH + length_pulse
102
103             lines = [[(0.0, 0.0), (TOTAL_WIDTH, 0.0)]]
104             for spawn_x in GLOBAL_SPAN_POINTS:
105                 grow_branch(spawn_x, 0, 90, current_length, current_angle, MAX_DEPTH, lines)
106
107             mls = MultiLineString(lines)
108
109             # Original buffer: full nozzle radius, fully rounded joins/caps. No closing, no simplify.
110             poly = mls.buffer(NOZZLE_DIA / 2.0, join_style=1, cap_style=1)
111
112             if poly.geom_type == "MultiPolygon":
113                 poly = max(poly.geoms, key=lambda a: a.area)
114
115             GLOBAL_POLYGONS[global_layer_idx] = poly
116
117         def corrugated_back_pts(x_start, x_end):
118             dist = abs(x_end - x_start)
119             n = max(4, int(dist / BACK_CORRUGATION_WAVE * BACK_CORRUGATION_PTS))
120             return []
121             (x_start + (x_end - x_start) * i / n,
122              Y_BACK + BACK_CORRUGATION_AMP * math.sin(2 * math.pi * (x_start + (x_end - x_start) * i / n) / BACK_CORRUGATION_WAVE))
123             for i in range(n + 1)
124
125         def generate_layer_perimeter(layer_idx, row, col):
126             global_layer_idx = row * NUM_LAYERS + layer_idx
127             layer_z = layer_idx * LAYER_HEIGHT

```

```

128         global_poly = GLOBAL_POLYGONS[global_layer_idx]
129
130         X_MIN_GLOBAL = col * TILE_WIDTH + X_L
131         X_MAX_GLOBAL = col * TILE_WIDTH + X_R
132
133         tile_clip_box = Polygon([
134             (X_MIN_GLOBAL, -100.0),
135             (X_MAX_GLOBAL, -100.0),
136             (X_MAX_GLOBAL, Y_MAX + 100.0),
137             (X_MIN_GLOBAL, Y_MAX + 100.0)
138         ])
139
140         local_poly = global_poly.intersection(tile_clip_box)
141         if local_poly.geom_type == "MultiPolygon":
142             local_poly = max(local_poly.geoms, key=lambda a: a.area)
143
144         shifted_coords = [(p[0] - col * TILE_WIDTH, p[1]) for p in local_poly.exterior.coords]
145
146         left_target = (X_L, 0.0)
147         right_target = (X_R, 0.0)
148
149         reordered = reorder_path(shifted_coords, left_target)
150         right_idx = min(range(len(reordered)), key=lambda i: distance(reordered[i], right_target))
151
152         path_a = reordered[:right_idx+1]
153         path_b = reordered[right_idx:]
154         avg_y = lambda pts: sum(p[1] for p in pts) / len(pts) if pts else 0
155         wavy_front = path_a if avg_y(path_a) > avg_y(path_b) else path_b[:-1]
156
157         if wavy_front and wavy_front[0][0] > wavy_front[-1][0]:
158             wavy_front = wavy_front[::-1]
159
160         wavy_front = [p for p in wavy_front if p[1] >= Y_WAVE_BASE - 0.5]
161         if wavy_front:
162             wavy_front[0] = (wavy_front[0][0], Y_WAVE_BASE)
163             wavy_front[-1] = (wavy_front[-1][0], Y_WAVE_BASE)
164
165         EYELET_X = TILE_WIDTH * 0.25 if row % 2 != 0 else TILE_WIDTH * 0.75
166         EYELET_Y = Y_BACK + EYELET_INSET
167         hw = EYELET_NECK / 2.0; r = EYELET_DIAMETER / 2.0
168         if hw == r: hw = r - 0.1
169         dy = math.sqrt(r**2 - hw**2)
170         a1 = math.atan2(-dy, hw); a2 = math.atan2(-dy, -hw) + 2*math.pi
171         arc_pts = [(EYELET_X + r*math.cos(a1 + (a2-a1)*(i/24.0)),
172                    EYELET_Y + r*math.sin(a1 + (a2-a1)*(i/24.0))] for i in range(25)]
173
174         lp = list(wavy_front)
175         lp += [(X_R, Y_WAVE_BASE), (X_R, Y_NOTCH_TOP), (X_R-NOTCH_INSET, Y_NOTCH_TOP),
176              (X_R-NOTCH_INSET, Y_NOTCH_BOTTOM), (X_R, Y_NOTCH_BOTTOM)]
177         if layer_z <= EYELET_Z_HEIGHT:
178             lp += corrugated_back_pts(X_R, EYELET_X+hw)
179             lp += [(EYELET_X+hw, EYELET_Y-dy)]
180         else:
181             lp += corrugated_back_pts(X_R, X_L)
182             lp += [(X_L, Y_NOTCH_BOTTOM), (X_L-NOTCH_INSET, Y_NOTCH_BOTTOM),
183                  (X_L-NOTCH_INSET, Y_NOTCH_TOP), (X_L, Y_NOTCH_TOP),
184                  (X_L, Y_WAVE_BASE), wavy_front[0]]
185
186         clean = []
187         for pt in lp:
188             if not clean or distance(clean[-1], pt) > 0.01:
189                 clean.append(pt)
190         return clean
191
192         # 5. DATA STRUCTURES
193         # =====
194         all_tiles_gcode = []
195
196         # 6. TILE GENERATOR
197         # =====
198         def generate_tile(row, col):
199             filename = f"Tile_R{row}_C{col}.gcode"
200             print(f"Generating {filename}...")
201             tile_gcode_lines = []
202
203             def wg(line, ext=True): tile_gcode_lines.append(line, ext)
204
205             wg(f"Biophilic Cladding - Row {row}, Column {col}\n", False)
206             wg("G90\nM82\nG28\nG92\n", False)
207             cur_e = 0.0; cur_z = 0.0
208
209             base_perimeter = generate_layer_perimeter(0, row, col)
210             if not base_perimeter:
211                 return
212             base_polygon = Polygon(base_perimeter)
213             step = NOZZLE_DIA * 0.75
214             current_poly = base_polygon.buffer(-NOZZLE_DIA * 0.6)
215             solid_path = list(base_perimeter)
216
217             while not current_poly.is_empty():
218                 if current_poly.geom_type == "MultiPolygon":
219                     current_poly = max(current_poly.geoms, key=lambda a: a.area)
220                 elif current_poly.geom_type == "Polygon":
221                     ring_coors = list(current_poly.exterior.coords)
222                     reorder_coors = reorder_path(ring_coors, solid_path[-1])
223                     solid_path.extend(ring_coors)
224                     current_poly = current_poly.buffer(-step)
225
226                 for layer in range(BOTTOM_SOLID_LAYERS):
227                     cur_z += LAYER_HEIGHT
228                     wg(f"\n; --- SOLID BOTTOM LAYER {layer} ---\n", False)
229                     wg(f"G1 Z{cur_z:.4f} F{Z_FEEDRATE}\n", False)
230                     sp = solid_path[0]
231                     wg(f"G0 X{sp[0]:.3f} Y{sp[1]:.3f} F{XY_TRAVEL_FEEDRATE}\n", False)
232                     for i in range(1, len(solid_path)):
233                         pt = solid_path[i]
234                         d = distance(solid_path[i-1], pt)
235                         cur_e += d * E_MULTIPLIER
236                     wg(f"G1 X{pt[0]:.3f} Y{pt[1]:.3f} E{cur_e:.5f} F{FEEDRATE}\n")
237
238                 for layer in range(BOTTOM_SOLID_LAYERS, NUM_LAYERS):
239                     wg(f"\n; --- VASE MODE LAYER {layer} ---\n", False)
240                     path = generate_layer_perimeter(layer, row, col)
241                     if not path: continue
242                     total_len = sum(distance(path[i-1], path[i]) for i in range(1, len(path)))
243                     if total_len == 0: continue
244                     sp = path[0]
245                     bridge = distance(solid_path[-1], sp) if layer == BOTTOM_SOLID_LAYERS else 0
246                     cur_e += bridge * E_MULTIPLIER
247                     wg(f"G1 X{sp[0]:.3f} Y{sp[1]:.3f} E{cur_e:.5f} F{FEEDRATE}\n")
248                     for i in range(1, len(path)):
249                         pt = path[i]; d = distance(path[i-1], pt)

```

```

249                         cur_e += d * E_MULTIPLIER
250                     wg(f"G1 X{pt[0]:.3f} Y{pt[1]:.3f} Z{cur_z:.4f} E{cur_e:.5f} F{FEEDRATE}\n")
251
252                     cur_z = (d / total_len) * LAYER_HEIGHT; cur_e += d * E_MULTIPLIER
253                     wg(f"G1 X{pt[0]:.3f} Y{pt[1]:.3f} Z{cur_z:.4f} E{cur_e:.5f} F{FEEDRATE}\n")
254
255                     with open(filename, 'w', encoding='utf-8') as f:
256                         for line, _ in tile_gcode_lines: f.write(line)
257
258                     # --- FIX: Append generated data to the global list for the HTML preview ---
259                     all_tiles_gcode.append({
260                         'row': row,
261                         'col': col,
262                         'lines': tile_gcode_lines
263                     })
264
265                 # 7. RUN
266                 # =====
267                 print(f"Processing Complete.")
268
269                 precompute_global_polygons()
270
271                 for r in range(MATRIX_ROWS):
272                     for c in range(MATRIX_COLS):
273                         generate_tile(r, c)
274
275                 print(f"Structuring Complete.")
276
277                 # 8. MASTER PREVIEW G-CODE & JSON STRUCTURING
278                 # =====
279                 frontend_data = {
280                     "rows": MATRIX_ROWS,
281                     "cols": MATRIX_COLS,
282                     "tileWidth": TILE_WIDTH,
283                     "zMax": Z_MAX,
284                     "tiles": []
285                 }
286                 for td in all_tiles_gcode:
287                     r, c = td['row'], td['col']
288                     x_off = c * TILE_WIDTH
289                     z_off = r * Z_MAX
290                     raw_gcode = "".join([line for line, _ in td['lines']])
291                     tile_segments = []
292                     cx = cy = cz = 0.0
293                     for line, is_ext in td['lines']:
294                         if line.startswith("G0", "G1"):
295                             parts = line.split()
296                             nx, ny, nz = cx, cy, cz
297                             for p in parts[1:]:
298                                 if p.startswith("X"): nx = float(p[1:]) + x_off
299                                 elif p.startswith("Y"): ny = float(p[1:])
300                                 elif p.startswith("Z"): nz = float(p[1:]) + z_off
301
302                                 seg = math.sqrt((nx-cx)**2 + (ny-cy)**2 + (nz-cz)**2)
303                                 if is_ext and seg > 0.01:
304                                     tile_segments.extend([round(cx,1), round(cy,1), round(cz,1),
305                                                            round(nx,1), round(ny,1), round(nz,1)])
306
307                                 cx, cy, cz = nx, ny, nz
308
309                 frontend_data["tiles"].append({
310                     "row": r,
311                     "col": c,
312                     "gcode": raw_gcode,
313                     "segments": tile_segments
314                 })
315
316                 # 9. INTERACTIVE HTML PREVIEW
317                 # =====
318                 HTML_TEMPL = """
319                 <DOCTYPE html>
320                 <html lang="en">
321                 <head>
322                     <meta charset="utf-8">
323                     <title>Terracotta Biophilic Facade UI</title>
324                     <style>
325                         *{margin:0;padding:0;box-sizing:border-box}
326                         body{overflow:hidden;background:#f0f4f8;font-family:'Courier New',monospace}
327
328                         #sidebar-content{font-size:13px;color:#444;margin-bottom:20px;}
329                         .btn{display:block;width:100%;padding:10px;background:#1a7388;color:white;text-align:center;border:none;}
330                         .btn:hover{background:#1557b0;}
331                         .btn:disabled{background:#ddd;color:#888;cursor:not-allowed;}
332
333                     </style>
334                 </head>
335                 <body>
336                     <div id="loading">Building Interactive Matrix</div>
337                     <div id="ui">
338                         <h1>Biophilic Facade</h1>
339                         <p id="stat"></p>
340                         <p style="margin-top:8px">Left-Click Brick to Select &mdot; Drag to Rotate &mdot; Right-Drag to Pan</p>
341
342                         <div id="sidebar">
343                             <h2>Inspector</h2>
344                             <div id="sidebar-content">Click on a brick in the 3D view to inspect it and download its toolpath.</div>
345                             <button id="download-btn" class="btn" disabled=Download G-Code</button>
346                         </div>
347
348                         <script src="https://cdnjs.cloudflare.com/ajax/libs/three.js/r128/three.min.js"></script>
349                         </script>
350                         /VARS/
351
352                         const ZT = FACADE_DATA.rows * FACADE_DATA.zMax;
353                         const XT = FACADE_DATA.cols * FACADE_DATA.tileWidth;
354
355                         document.getElementById('stat').innerHTML =

```

```

356         '<span class="tag">' + FACADE_DATA.rows + '&times;' + FACADE_DATA.cols + '</span> Matrix';
357
358         const scene = new THREE.Scene();
359         scene.background = new THREE.Color(0xf0f4f8);
360         scene.fog = new THREE.Fog(0xf0f4f8, 1000, 5000);
361
362         const cam = new THREE.PerspectiveCamera(40, innerWidth/innerHeight, 1, 8000);
363         const renderer = new THREE.WebGLRenderer({antialias:true, powerPreference:'high-performance'});
364         renderer.setPixelRatio(Math.min(devicePixelRatio, 2));
365         renderer.setSize(innerWidth, innerHeight);
366         renderer.shadowMap.enabled = true;
367         renderer.shadowMap.type = THREE.PCFSoftShadowMap;
368         renderer.toneMapping = THREE.ACESFilmicToneMapping;
369         renderer.toneMappingExposure = 1.0;
370         document.body.appendChild(renderer.domElement);
371
372         scene.add(new THREE.AmbientLight(0xfffff, 0.6));
373
374         const sun = new THREE.DirectionalLight(0xffff50, 0.8);
375         sun.position.set(XT/2, ZT+1.5, 1000);
376         sun.castShadow = true;
377         scene.add(sun);
378
379         const fillFront = new THREE.DirectionalLight(0x44444, 0.4);
380         fillFront.position.set(-XT, ZT/2, 1200);
381         scene.add(fillFront);
382
383         const gnd = new THREE.Mesh(
384             new THREE.PlaneGeometry(10000, 10000),
385             new THREE.MeshStandardMaterial({color:0x8e8e8e, roughness:0.9})
386         );
387         gnd.rotation.x = -Math.PI/2;
388         gnd.position.set(XT/2, -1, 0);
389         gnd.receiveShadow = true;
390         scene.add(gnd);
391
392         const cyGeo = new THREE.CylinderGeometry(NR, NR, 1.0, 8, 1);
393
394         const cyMat = new THREE.MeshStandardMaterial({
395             roughness: 0.3,
396             metalness: 0.1,
397             color: 0xfffff
398         });
399         const dummy = new THREE.Object3D();
400         const UP = new THREE.Vector3(0, 1, 0);
401         const dirV = new THREE.Vector3();
402
403         const tileMeshes = [];
404         const hitboxes = [];
405         const hitboxGeo = new THREE.BoxGeometry(FACADE_DATA.tileWidth, FACADE_DATA.zMax, 60);
406         const hitboxMat = new THREE.MeshBasicMaterial({visible: false});
407
408         FACADE_DATA.tiles.forEach(tile => {
409             const segs = tile.segments;
410             const numSegs = segs.length / 6;
411             if(numSegs === 0) return;
412
413             // --- Calculate Relative Tips using Local Moving Average ---
414             const ridgeVals = new Float32Array(numSegs);
415
416             // WIDER window to establish a true baseline for the curve
417             const WINDOW = 35;
418
419             for(let i=0; i<numSegs; i++) {
420                 const o_curr = i * 6;
421                 const curr_y = (segs[o_curr+1] + segs[o_curr+4]) * 0.5; // Python Y (depth)
422
423                 // Look ONLY at points safely on the front face (past the 5mm mark)
424                 if (curr_y > 5.0) {
425                     let sumY = 0;
426                     let count = 0;
427
428                     for(let j = -WINDOW; j <= WINDOW; j++) {
429                         const idx = i + j;
430                         if(idx >= 0 && idx < numSegs) {
431                             const o_adj = idx * 6;
432                             const adj_y = (segs[o_adj+1] + segs[o_adj+4]) * 0.5;
433
434                             // CRITICAL FIX: Only include other front-face points in the average!
435                             // This prevents the deep back-notches from dragging down the average
436                             // and creating false "peaks" at the seams of the tiles.
437                             if (adj_y > 5.0) {
438                                 sumY += adj_y;
439                                 count++;
440                             }
441                         }
442                     }
443
444                     if (count > 0) {
445                         const avgY = sumY / count;
446                         ridgeVals[i] = Math.max(0, curr_y - avgY);
447                     }
448                 }
449             }
450
451             const tileMat = cyMat.clone();
452             tileMat.transparent = true;
453
454             const inst = new THREE.InstancedMesh(cyGeo, tileMat, numSegs);
455             inst.castShadow = inst.receiveShadow = true;
456
457             const baseColor = new THREE.Color(0x556b2f); // Moss Green
458             const tipColor = new THREE.Color(0xf4f4f0); // Off White
459
460             for (let i = 0; i < numSegs; i++) {
461                 const o = i * 6;
462                 const ax = segs[o+2], ay = segs[o+2], az = -segs[o+1];
463                 const bx = segs[o+3], by = segs[o+3], bz = -segs[o+4];
464                 const dx=bx-ax, dy=by-ay, dz=bz-az;
465                 const len = Math.sqrt(dx*dx+dy*dy+dz*dz);
466
467                 if (len < 0.1) { inst.setMatrixAt(i, new THREE.Matrix4().makeScale(0,0,0)); continue; }
468
469                 dummy.position.set((ax+bx)*0.5, (ay+by)*0.5, (az+bz)*0.5);
470                 dummy.scale.set(1, len, 1);
471                 dirV.set(dx/len, dy/len, dz/len);
472                 if (Math.abs(dirV.y) > 0.999) dummy.quaternion.set(0, 0, dirV.y < 0 ? 1 : 0, dirV.y < 0 ? 0 : 1);
473                 else dummy.quaternion.setFromUnitVectors(UP, dirV);
474                 dummy.updateMatrix();
475                 inst.setMatrixAt(i, dummy.matrix);
476
477                 const clr = new THREE.Color();
478                 clr.copy(baseColor);
479
480                 if (ridgeVals[i] > 0) {
481                     // Map the depth difference to an intensity scale.
482                     // Divided by 2.5 to make the blue slightly more prominent on true peaks

```

final_corg.py (continued)

```

513 let intensity = Math.min(1.0, ridgeVals[i] / 2.5);
514
515 // Smoothstep function for butter-smooth color blending
516 intensity = intensity * intensity * (3 - 2 * intensity);
517
518 clr.lerp(tipColor, intensity);
519
520
521 inst.setColorAt(1, clr);
522
523
524 scene.add(inst);
525 tileMeshes.push(inst);
526
527 const box = new THREE.Mesh(hitboxGeo, hitboxMat);
528 box.position.set(
529   tile.col * FACADE_DATA.tileWidth + (FACADE_DATA.tileWidth / 2),
530   tile.row * FACADE_DATA.zMax + (FACADE_DATA.zMax / 2),
531   -10
532 );
533 box.userData = { tileInfo: tile, meshIndex: tileMeshes.length - 1 };
534 scene.add(box);
535 hitboxes.push(box);
536 });
537
538 let theta = -0.6, phi = 1.2, dist = 1200;
539 const target = new THREE.Vector3(XT/2, ZT*0.4, 0);
540
541 function syncCam() {
542   cam.position.set(
543     target.x + dist * Math.sin(phi) * Math.sin(theta),
544     target.y + dist * Math.cos(phi),
545     target.z + dist * Math.sin(phi) * Math.cos(theta)
546   );
547   cam.lookAt(target);
548 }
549 syncCam();
550
551 const raycaster = new THREE.Raycaster();
552 const mouse = new THREE.Vector2();
553 let selectedTileData = null;
554
555 let drag=false, pan=false, lx=0, ly=0, moved=false;
556
557 window.addEventListener('mousedown', e => {
558   if(e.target.tagName === 'BUTTON') return;
559   if(e.button===0) drag=true;
560   if(e.button===2) pan=true;
561   lx=e.clientX; ly=e.clientY;
562   moved=false;
563 });
564
565 window.addEventListener('mouseup', e => {
566   drag=pan=false;
567   if(!moved && e.button === 0 && e.target.tagName !== 'BUTTON') {
568     handleRaycast(e);
569   }
570 });
571
572 window.addEventListener('mousemove', e => {
573   if (drag) {
574     theta += (e.clientX-lx)*0.005;
575     phi = Math.max(0.1, Math.min(Math.PI-0.1, phi-(e.clientY-ly)*0.005));
576     syncCam();
577     moved=true;
578   }
579   if (pan) {
580     const right = new THREE.Vector3().crossVectors(new THREE.Vector3().subVectors(target, cam.position), new THREE.Vector3(0,1,0).normalize());
581     target.addScaledVector(right, -(e.clientX-lx)*dist*0.001);
582     target.y += (e.clientY-ly)*dist*0.001;
583     syncCam();
584     moved=true;
585   }
586   lx=e.clientX; ly=e.clientY;
587 });
588 window.addEventListener('wheel', e => { dist = Math.max(100, Math.min(4000, dist * (1+e.deltaY*0.001))); syncCam(); }, {passive:true});
589
590 function handleRaycast(event) {
591   mouse.x = (event.clientX / window.innerWidth) * 2 - 1;
592   mouse.y = -(event.clientY / window.innerHeight) * 2 + 1;
593
594   raycaster.setFromCamera(mouse, cam);
595   const intersects = raycaster.intersectObjects(hitboxes);
596
597   if (intersects.length > 0) {
598     const hitData = intersects[0].object.userData;
599     selectedTileData = hitData.tileInfo;
600
601     tileMeshes.forEach((mesh, index) => {
602       if(index === hitData.meshIndex) {
603         mesh.material.opacity = 1.0;
604       } else {
605         mesh.material.opacity = 0.2;
606       }
607     });
608
609     document.getElementById('sidebar-content').innerHTML = `
610 <strong>Selected Tile:</strong> Row ${selectedTileData.row}, Column ${selectedTileData.col} <br><br>
611 <strong>Dimensions:</strong> ${FACADE_DATA.tileWidth}mm x ${FACADE_DATA.zMax}mm <br>
612 <strong>Lines of G-code:</strong> ${selectedTileData.gcode.split('\n').length.toLocaleString()}
613 `;
614     const btn = document.getElementById('download-btn');
615     btn.disabled = false;
616     btn.onclick = () => {
617       const blob = new Blob([selectedTileData.gcode], { type: 'text/plain' });
618       const url = window.URL.createObjectURL(blob);
619       const a = document.createElement('a');
620       a.href = url;
621       a.download = `TILE_${selectedTileData.row}_C${selectedTileData.col}.gcode`;
622       document.body.appendChild(a);
623       a.click();
624       document.body.removeChild(a);
625       window.URL.revokeObjectURL(url);
626     };
627   } else {
628     selectedTileData = null;
629     tileMeshes.forEach(mesh => mesh.material.opacity = 1.0);
630     document.getElementById('sidebar-content').innerHTML = "Click on a brick in the 3D view to inspect it and download its toolpath.";
631     document.getElementById('download-btn').disabled = true;
632   }
633 }
634
635 const loadDiv = document.getElementById('loading');
636 (function loop(){
637   requestAnimationFrame(loop);
638   renderer.render(scene, cam);
639   if(loadDiv) { loadDiv.style.opacity='0'; setTimeout(()=>loadDiv.remove(),600); }
640 })();

```

```

641 </script>
642 </body>
643 </html>
644 ---
645
646 html_out = HTML_TEMPL.replace("/{VARS}/", JS_VARS)
647
648 print("\nGenerating Files...")
649 with open("FACADE_PREVIEW.html", 'w', encoding='utf-8') as f:
650   f.write(html_out)
651
652 print("--*40)
653 print("SUCCESS: INTERACTIVE PREVIEW GENERATED")
654 print("Check FACADE_PREVIEW.html")
655 print("--*40)

```

Thermal measurements (infrared)

FLIR surface-temperature readings, 22–23 June. Baselines are bare urban surfaces; experimental readings are the bioreceptive tile (moss and bare-ceramic columns).

Baseline — bare urban surfaces

Time	File	°C	Surface / placement
22 Jun 11:35	FLIR0244	41.8	wooden panel
22 Jun 11:35	FLIR0245	37.4	wooden panel
22 Jun 13:00	FLIR0255	48.1	decking
22 Jun 20:31	FLIR0269	46.9	brick facade
23 Jun 08:31	FLIR0273	25.3	stucco wall

Experimental — bioreceptive tile (TEST 2)

Time	File	°C	Surface / placement
22 Jun 11:33	FLIR0236	29.5	center · report image
22 Jun 11:33	FLIR0237	27.2	center · report image
22 Jun 11:33	FLIR0238	29.6	moss tile middle
22 Jun 11:33	FLIR0239	25.6	no moss tile middle
22 Jun 11:33	FLIR0240	30.4	moss tile middle
22 Jun 11:33	FLIR0241	25.6	no moss middle
22 Jun 11:33	FLIR0242	30.5	moss middle
22 Jun 11:33	FLIR0243	25.6	no moss center
22 Jun 12:58	FLIR0246	29.3	top no moss
22 Jun 12:58	FLIR0247	30.5	middle no moss
22 Jun 12:58	FLIR0248	34.8	bottom no moss
22 Jun 12:58	FLIR0250	33	bottom moss
22 Jun 12:58	FLIR0251	35.5	middle moss
22 Jun 12:58	FLIR0252	36.1	middle moss
22 Jun 12:58	FLIR0253	27.9	top moss
22 Jun 12:58	FLIR0254	33.6	top moss
22 Jun 12:58	FLIR0258	34.5	middle moss
22 Jun 12:58	FLIR0260	32.7	bottom moss
22 Jun 20:30	FLIR0261	25.3	middle no moss
22 Jun 20:30	FLIR0262	25.3	middle moss
22 Jun 20:30	FLIR0263	23.4	top moss
22 Jun 20:30	FLIR0264	23.4	top no moss
22 Jun 20:30	FLIR0265	25.9	bottom no moss
22 Jun 20:30	FLIR0267	24	bottom moss
23 Jun 08:32	FLIR0270	19.8	middle moss
23 Jun 08:32	FLIR0271	19.9	middle no moss

Ambient air — KNMI station 344 (Rotterdam Airport)

Hourly air temperature, converted to local summer time. Nearest official station to Delft (~7 km).

22 Jun, late morning (~11:35)	23.5 °C
22 Jun, midday (~13:00)	25.5 °C
22 Jun, afternoon peak (15:00–17:00)	26.3 °C
22 Jun, evening (~20:30)	24.5 °C
Overnight minimum	~16.5 °C
23 Jun, next morning (~08:30)	19.0 °C

Source: KNMI hourly data (uurgegevens), station 344.
 daggegevens.knmi.nl/klimatologie/uurgevens

On-site reading at 19:58 on 22 Jun was 26 °C, consistent with the KNMI value for a sheltered urban courtyard.