



priukrasit



editor-in-chief rhea agrawal
the bulleteen

from the editor's desk

Fashion doesn't usually get filed under "STEM." It gets filed under art, or business, or maybe culture. But underneath almost every interesting textile, dye, or garment is a real materials science, chemistry, or computation problem that someone had to solve. These eight pieces each start from a piece of clothing and end up somewhere more technical than you'd expect.


Because it is through the way we adorn ourselves that we express the different ways we engage with the world around us and show parts and pieces of ourselves hidden deep within facades and performance. “Priukrasit,” which means to adorn or dress up in Russian, is an ode to this idea: that there is something deeply scientific about the ways in which we choose to present ourselves to the world.

A dress can be a chemical problem. A piece of jewellery can become a question of materials and structure. A printed pattern can involve mathematics, geometry, and computation. Even something as familiar as the colour of a shirt carries with it an entire chain of science, from the chemistry of pigments and dyes to the physics of how light interacts with a surface. What we see as decoration is often the visible result of an invisible process of experimentation, engineering, and problem-solving.

And perhaps that is what makes fashion such an interesting way into STEM. Science is often imagined as something distant from the ordinary: equations on a page, machines in a laboratory, data on a screen. Fashion exists in the opposite space. It is intimate, familiar, and human. We wear it, touch it, alter it, inherit it, and use it to tell stories about ourselves. Looking at fashion through STEM makes those invisible scientific processes visible, while also showing that science does not exist separately from the ways we live and create.

“Priukrasit” is therefore not simply about dressing up. It is about understanding the science hidden within the things we use to dress ourselves up. It asks us to look past the finished garment and wonder what had to be understood, engineered, calculated, or discovered for it to exist in the first place.

The eight pieces that follow are small explorations of that intersection. Each begins with something we might normally see as fashion, but follows it somewhere unexpected: into chemistry, physics, biology, materials science, mathematics, or computation. Together, they suggest that perhaps STEM has always been closer to fashion than we thought.



Yours Truly,
Rhea Agrawal
Founder and Director of the Bulleteen

table of contents



01

THE PHYSICS BEHIND A BUTTERFLY WING (AND WHY IT NEVER FADES)

By Teela Steinert

02

WHY IT'S SO HARD TO PUT A CIRCUIT IN A WASHING MACHINE?

By Fathima Zahra

03

CAN AN ALGORITHM ACTUALLY DESIGN A DRESS?

By Aamnah Younus

04

GROWING LEATHER INSTEAD OF MAKING IT

By Alexa Kightlinger

05

WHEN IS A GARMENT ACTUALLY A MACHINE?

By Danielle Patton

06

THE MATERIALS THAT REMEMBER THEIR SHAPE

By Chang Lu

07

THE CHEMISTRY NOBODY THINKS ABOUT WHEN THEY BUY A T-SHIRT

By Traicy Moren

08

WHAT SHARKSKIN AND LOTUS LEAVES TAUGHT TEXTILE ENGINEERS

By Feroqh Shakouri

THE PHYSICS BEHIND A BUTTERFLY WING (AND WHY IT NEVER FADES)

By Teela Steinert

Grind a piece of blue morpho butterfly wing into a fine enough powder, and something strange happens: the color disappears. What's left is a dull brown-gray dust. That's the tell that the morpho's famous, almost impossibly saturated blue isn't a pigment at all. There's no blue chemical compound sitting in those wing scales waiting to be dissolved or worn away. The color is structural: it's produced entirely by the physical architecture of the wing scale interacting with light, and once that architecture is destroyed, the color goes with it.



Here's the distinction that matters. Pigment-based color works by selective absorption: a red pigment absorbs every wavelength of visible light except red, which bounces back to your eye. Structural color works by interference and diffraction instead. Morpho wing scales are covered in stacked, ridged layers of a biological material called chitin, arranged at a periodicity of a few hundred nanometers, close to the wavelength of visible light itself. When light hits these layers, different wavelengths reflect off different depths and interfere with each other, some constructively and some destructively, and the wavelengths that survive that interference are what you perceive as color. Tilt the wing and the angle of light striking those layers changes, which shifts which wavelengths interfere constructively. That's why the morpho's blue seems to shimmer and shift as it flies, an effect pigment can't produce.


This same trick shows up across the natural world in slightly different geometries: peacock feathers use a photonic-crystal-like lattice of melanin rods, opals use a regular stack of silica spheres, and some beetle shells reflect circularly polarized light, an effect unusual enough that it's still an active area of biophysics research. What unites all of these is that the color comes from nanoscale periodicity, not chemistry.

You've seen a milder version of the same phenomenon without ever looking at an insect: the rainbow sheen on a soap bubble, or a thin film of oil on wet pavement, works through the same basic mechanism, thin-film interference between light reflecting off the top and bottom of an extremely thin layer. The difference is that a butterfly wing's structure is repeated precisely and consistently across billions of nanostructures on a single wing, rather than existing as one uneven, shifting film, which is why the resulting color is vivid and stable rather than the fleeting, unstable rainbow you get off a puddle.

For textile designers, this is attractive for a fairly practical reason: pigments fade. UV exposure breaks down the chemical bonds that make a dye molecule absorb the wavelengths it does, which is why a red shirt left in direct sun eventually looks pink. A structural color, in principle, doesn't fade the same way, because there's no chemical bond being broken. The color is a geometric property of the surface, and as long as the geometry survives, so does the color. Structural colorants would also sidestep a lot of the toxic chemistry involved in synthetic dye production, since you're manipulating physical structure rather than synthesizing and fixing a pigment molecule onto a fiber.

The challenge is manufacturing that periodicity at any usable scale. A butterfly wing achieves its nanostructure through a biological self-assembly process refined over evolutionary time. Replicating a few-hundred-nanometer-precision structure across meters of woven or knitted fabric, cheaply enough to compete with a dye vat, is a genuinely hard materials engineering problem. Researchers working with cellulose nanocrystals, derived from plant fiber and capable of self-assembling into similarly periodic structures, have made real progress producing structurally colored films. Matching that durability and vibrancy on something as mechanically abused as an actual garment, though, is still mostly a research-stage problem rather than a store-shelf one.

There's a durability irony here, too: the same delicate nanostructure that gives you fade-proof color is often more fragile to physical wear than a dyed fiber is. A dye molecule is chemically bonded throughout the fiber; a structural color depends on a precise surface geometry that abrasion can flatten or scratch away. So the field's real open question isn't whether structural color can be made. Researchers already can, in the lab. It's whether that color can survive contact with an actual human wardrobe.






Illustration by Ralph McQuarrie for Isaac Asimov's 1986 book, Robot Dreams

WHY IT'S SO HARD TO PUT A CIRCUIT IN A WASHING MACHINE?

By Fathima Zahra

The hardest part of building a smart shirt was never getting a sensor to work. It's getting that sensor to still work after the shirt has been through a washing machine fifty times, stretched across someone's shoulders while they reach overhead, balled up in a gym bag, and generally treated the way people treat clothing rather than the way people treat electronics.

Getting conductivity into a fiber in the first place is, on its own, a fairly solved problem with a few different approaches. Metal-core yarns work by spinning or wrapping a fine metal filament, often silver-coated copper or stainless steel, together with conventional textile fiber like nylon or polyester, producing a yarn that's mechanically similar to normal thread but electrically conductive along its length. Conductive polymer coatings take a different approach, applying a coating like PEDOT:PSS (a polymer that happens to be both electrically conductive and reasonably flexible) directly onto conventional fiber, letting it carry current across its surface. Either approach can be woven or knitted using largely conventional textile machinery, which is part of why the sensing side of e-textiles has moved quickly.

The durability problem shows up once that conductive element has to survive being a garment rather than a lab sample. Fabric flexes constantly and unpredictably. A fiber running along someone's elbow gets bent thousands of times a day in ways a rigid circuit board never experiences, and repeated flexing at the same point causes mechanical fatigue that can crack a metal filament or wear through a polymer coating over time. Washing introduces a second layer of stress entirely: water, detergent chemistry, mechanical agitation, and heat from drying all attack conductive coatings and metal filaments in ways ordinary fiber shrugs off.




One partial fix, borrowed from stretchable electronics research more broadly, is to stop trying to make the conductor itself stretchy, and instead change its geometry. A straight length of conductive trace snaps quickly under repeated stretching, but the same trace laid out in a wavy, serpentine, or horseshoe-shaped path can absorb a surprising amount of physical elongation just by unfurling that waviness, distributing strain across the whole path rather than concentrating it at one point. The trace itself never needs to stretch as a material; it just needs room to unfold. Versions of this geometric approach, originally developed for flexible skin-mounted medical sensors, have started showing up in e-textile interconnect design for the same reason.

This is why the e-textile products that have made it to market tend to be narrower in scope than the sweeping "smart clothing" vision often suggests. Medical-grade smart shirts that monitor heart rate or respiration through embedded electrodes, for instance, tend to isolate the rigid, non-washable components (batteries, main processing units) into small removable modules that snap out before laundering, while only the flexible conductive fiber stays sewn into the fabric itself. That's a pragmatic compromise, not a fully solved problem. Fabric and rigid electronics are still, at a materials level, fundamentally different substances with different failure modes, and the most durable e-textile designs are the ones that stop trying to force full integration and instead design around that mismatch.

CAN AN ALGORITHM ACTUALLY DESIGN A DRESS?

By Aamnah Younus



Traditional pattern-making is a deeply manual craft: a designer drapes fabric on a form, or draws flat pattern pieces by hand, iterating through fittings and adjustments guided mostly by intuition and experience. Generative design flips that process around entirely. Instead of a designer directly specifying a shape, they specify constraints and goals (fabric width, seam placement rules, target measurements, desired silhouette parameters) and let a computational system search through the enormous space of possible solutions that satisfy those constraints, often producing pattern geometries a human designer wouldn't have arrived at through conventional intuition.

This works particularly well when paired with 3D body scanning. Instead of designing for a set of standardized sizes that approximate a population of bodies, a scanner captures an individual's actual measurements as a detailed point cloud, and generative software can adjust a garment's underlying pattern algorithmically to fit that specific body, rather than a designer manually grading a pattern up or down between fixed sizes. Companies working on on-demand and personalized apparel manufacturing have leaned on this approach specifically to reduce the overproduction and unsold inventory that comes from manufacturing standardized sizes speculatively, ahead of actual demand.

Generative design as a broader concept isn't unique to fashion. It's been used for years in fields like aerospace and automotive engineering, where software given a set of structural loads, material constraints, and manufacturing limits generates part geometries, often strange, organic-looking lattice shapes, that a human engineer wouldn't have

sketched but that satisfy the physical requirements more efficiently than a conventionally designed part. Applying the same approach to garment construction imports that same optimization logic into a domain where the "load" being satisfied is aesthetic and ergonomic rather than purely structural. That's part of why translating it well is harder than it sounds: fit and drape are far fuzzier constraints to encode mathematically than the load-bearing requirements of an airplane bracket.

A more provocative application involves generative adversarial networks, or GANs: a machine learning architecture with two competing neural networks, one (the generator) trying to produce convincing new outputs, and the other (the discriminator) trying to distinguish those generated outputs from real examples. Trained on large datasets of existing garment images or silhouettes, a GAN can be pushed to generate entirely novel garment designs, exploring the mathematical space between and beyond the examples it was trained on. Fashion-focused GAN research has produced generated silhouettes, patterns, and even full garment renderings that no human directly sketched, raising an interesting question about where creative authorship sits in that process.

That question deserves more than a shrug. What a GAN is doing, mathematically, is learning a compressed representation of its training data, a "latent space," and generating new outputs by sampling points within or near that space that weren't part of the original training examples. Interpolating between two points in that latent

space produces a "design" that blends the visual and structural features of the training examples it sits between. This differs from a designer's creative process in an important way: the network has no intention, no cultural context, no lived understanding of why a certain silhouette might communicate something to a wearer. It's optimizing a mathematical objective, usually producing outputs the discriminator network can't distinguish from real examples, not pursuing an expressive goal.

Whether that constitutes "design" depends on how strictly you define the word. If design requires intentionality, a purpose behind a choice, then a GAN isn't designing the way a human is; it's exploring a statistical space defined by its training data. But if design is treated more functionally, as the process of generating novel solutions that satisfy a set of constraints and produce a specific output, then generative systems arguably qualify, even without anything resembling intention behind the process. This is a version of a much older question about creativity and computation, one that predates fashion and shows up anywhere generative systems produce results people find surprising or aesthetically compelling.

Whatever you conclude about the philosophy, one thing is clear in practice: generative and computational pattern-making tools are changing who does the early-stage exploratory work in garment design. Some of it is shifting from a human iterating through physical drapes and sketches to a system iterating through a mathematically defined possibility space, with a human designer curating and selecting from what that system proposes rather than generating every option alone.



Detail of a garment from the Escapism collection (2011) by Iris van Herpen and Daniel Widrig



GROWING LEATHER INSTEAD OF MAKING IT

By Alexa Rightlinger

Somewhere between a farm and a factory, mycelium is quietly becoming one of the more promising materials in fashion's search for alternatives to animal leather and petroleum-based synthetics. Mycelium is the root-like network of fine white filaments that fungi grow through whatever substrate they're colonizing; the mushroom itself is just the fruiting body that pops up above ground. Grown under the right conditions, on a bed of agricultural waste like sawdust or corn stalks, that filament network forms a dense, felt-like mat within a matter of weeks, which can then be processed, tanned, and finished into a material with a leather-like hand-feel and drape.

Companies like MycoWorks and Bolt Threads have spent the past decade refining this into something closer to an industrial process than a science-fair project. The material has already made it into limited-run partnerships with major fashion houses: MycoWorks worked with Hermès on a mycelium-based bag prototype, and Bolt Threads' mycelium leather has appeared in collaborations with brands including Adidas and Stella McCartney. What makes this compelling from a sustainability standpoint isn't only that it avoids animal agriculture. The growth cycle is measured in weeks rather than years, and the inputs are largely agricultural byproducts rather than land, water, and feed.

The finishing process matters here almost as much as the growing process. A raw sheet of mycelium fresh out of its growing tray doesn't behave like leather at all: it's often too soft, too absorbent, or structurally inconsistent to use directly. Getting from that raw mat to something a shoe or bag manufacturer can cut, stitch, and stress-test the way they would a hide requires a separate finishing stage, compressing the material, treating its surface, sometimes reinforcing it with a backing layer. That finishing step is where most of the remaining engineering work in this field is happening now, since growing the raw material is, at this point, a comparatively well-understood process.

A parallel approach uses bacterial cellulose instead of fungal filament. If you've ever seen a kombucha SCOBY, the rubbery disc that forms on top of a batch of fermenting kombucha, you've seen bacterial cellulose being produced in real time. Bacteria in the fermenting liquid excrete cellulose fibers as a byproduct of their metabolism, and over days those fibers accumulate into a continuous sheet. Designer and researcher Suzanne Lee popularized growing and drying these sheets into wearable material in the 2000s under the name BioCouture, demonstrating that you could grow a jacket roughly the way you'd grow a batch of yogurt.




Biofabrication extends beyond leather substitutes into dyeing, too, and this might be the more quietly radical application. A company called Colorifix has engineered microorganisms to biosynthesize pigment directly, then apply it to fabric through a fermentation process rather than a chemical dye bath, growing the color onto the fiber instead of dyeing it on. This sidesteps a large portion of the water and heavy-metal pollution associated with conventional textile dyeing, because the process relies on a biological synthesis pathway instead of an industrial chemical one.

None of this is close to displacing conventional leather or synthetic dye at the scale the global fashion industry operates at, and the reasons are fairly straightforward. Growing a material biologically means being constrained by biological timeframes and yields. A bioreactor produces a finite batch over a finite number of days, and scaling that up isn't as simple as building a bigger vat, because the organisms need consistent conditions (temperature, humidity, nutrient supply) that get harder to maintain uniformly at industrial volume.

What makes biofabrication worth watching isn't that it's already won. It's a genuinely different manufacturing philosophy. Conventional leather and dye production is subtractive and additive in a mechanical, chemical sense: you strip, tan, and treat an animal hide, or you synthesize and fix a pigment molecule onto fiber. Biofabrication starts from a different question: instead of manufacturing a material, can you just grow the material you want, with the properties you want, from the start? It's a slower path today, but it scales with biology rather than against it, and biological processes have a long track record of eventually getting cheap.

WHEN IS A GARMENT ACTUALLY A MACHINE?

By Danielle Patton



Most robots, as commonly imagined, are built from rigid components: metal or hard plastic links connected by discrete joints, moved by motors, the same basic architecture whether you're picturing an industrial arm on an assembly line or a humanoid robot in a demo video. Soft robotics rejects that architecture almost entirely. Instead of rigid links and discrete joints, soft robots are built from compliant materials (silicone elastomers, flexible polymers, fabric) that deform continuously rather than bending at fixed pivot points, and are often actuated not by electric motors but by pneumatic or hydraulic pressure, inflating and deflating chambers embedded within the soft material to produce motion.

A common building block in soft robotic actuation is something called a pneu-net, short for pneumatic network: a molded elastomer structure containing a series of internal air chambers connected by a thin channel, engineered so that when pressurized, the chambers on one side expand more than a stiffer, less-expandable layer on the other side, causing the whole structure to curl or bend in a predictable direction. It's a purely geometric trick for turning uniform air pressure into directional motion, without needing any electric motor or moving mechanical part beyond the material itself flexing.

This design philosophy exists because rigid robots have a basic limitation around safety and adaptability: a rigid robotic arm moving with enough force to do useful work can also seriously injure a person it collides with, which is why industrial robots have traditionally been kept in caged, human-free work zones. A soft robotic system, built from materials with mechanical properties closer to human tissue, can safely make contact with a person by design, because its structure yields rather than impacts. That property, safe compliance rather than raw rigid force, is what makes soft robotics so well suited to wearable applications in a way rigid robotics isn't.


Exosuits are the clearest example of this crossing directly into apparel. Rather than the rigid, exoskeleton-style metal frameworks that show up in science fiction, research groups including Harvard's Wyss Institute have developed soft exosuits built almost entirely from textile components: fabric webbing, cable actuators, and pneumatic or cable-driven systems worn more like a harness than a suit of armor. These devices apply targeted mechanical assistance to specific muscle groups, timed to a wearer's own natural gait pattern, and have been developed for applications including gait rehabilitation for stroke patients and reducing fatigue for people doing physically demanding manual labor. Because the structure is fabric-based rather than rigid-metal-based, these devices can be worn under normal clothing and move with the body in a way a rigid exoskeleton structurally cannot.






This is where the categories start to blur. A soft exosuit is, depending on how you look at it, either a piece of clothing with embedded actuation or a robot built primarily from textile components, and which framing you use isn't just a semantic choice. It has real regulatory consequences. A device marketed and designed to treat a medical condition, like assisting gait for a stroke patient, generally has to go through a medical device approval pathway, with all the safety and efficacy testing that involves. A device marketed as general mobility or performance assistance for otherwise healthy people can sometimes avoid that same regulatory pathway entirely, even if the underlying soft-robotic technology is functionally similar or identical. That means the engineering itself (the actuators, the fabric, the control system) can be nearly the same across two products, while the regulatory category, and the testing burden that comes with it, is decided almost entirely by how the device is marketed and to whom.


This convergence raises a real definitional question without a clean answer: at what point does something stop being a garment with embedded mechanics and start being a robot you happen to wear? There's no obvious bright line. "Clothing" and "machine" were never separate categories with a hard boundary between them; they're points on a continuum of things that cover, support, and interact mechanically with a human body, and soft robotics is pushing textile engineering far enough along that continuum that the old categories stop applying cleanly. For soft wearable robotics specifically, fashion isn't a side concern. It may be one of the places this research has to succeed or fail, because the whole value proposition depends on a device being comfortable and wearable enough that someone actually chooses to put it on.



An Objects Tale

FIRST IMPRESSIONS LAST



A photograph of an astronaut floating in space, wearing a white spacesuit with red and blue accents. The astronaut is positioned diagonally across the frame, with their head towards the top right. The background is a deep blue and purple gradient, suggesting the Earth's atmosphere from space.

THE MATERIALS THAT REMEMBER THEIR SHAPE (AND REGULATE YOUR TEMPERATURE)

By Chang Lu

Astronauts have a temperature problem far more extreme than anything the rest of us deal with. Outside a spacesuit, the side of an object facing the sun can be baking while the shadowed side is freezing, sometimes by more than 200 degrees Celsius of difference. NASA's solution, developed in the 1980s, was to embed microscopic capsules of a wax-like material into spacesuit gloves: a material engineered to absorb heat by melting when things got too hot, and release that stored heat by resolidifying when things cooled down. That technology, licensed and commercialized under the name Outlast, is now the same basic principle behind a lot of "temperature-regulating" athletic wear sold today.

The physics behind it is called a phase-change material, or PCM, and the key concept is latent heat. When a substance changes phase (solid to liquid, for instance), it absorbs or releases a substantial amount of energy at a nearly constant temperature, energy that goes toward breaking or forming molecular bonds rather than raising or lowering temperature itself. That's why ice at 0°C takes a while to fully melt into water still at 0°C: the added heat energy is being consumed by the phase transition, not by warming the water further. PCM textiles exploit this directly. Microcapsules of a paraffin-based wax, engineered to melt near human skin temperature, are embedded into fiber or fabric coating. When your body starts to overheat, the surrounding capsules absorb that excess heat and melt, buffering your skin temperature. When you cool back down, the wax resolidifies and releases that stored heat back. The fabric isn't heating or cooling you in an active sense. It's acting as a thermal buffer, smoothing out temperature swings.

Shape-memory materials work on a completely different physical principle but toward a similarly adaptive goal. Shape-memory alloys, the best known being nitinol (a nickel-titanium alloy discovered at the Naval Ordnance Laboratory, which is where its name

comes from), can be deformed into a new shape at a lower temperature, then snap back to a pre-programmed original shape when heated past a specific transition temperature. This happens because of a reversible shift in the alloy's crystal structure between two phases: martensite, the more deformable low-temperature phase, and austenite, the more rigid high-temperature phase. Shape-memory polymers achieve a similar effect through a different mechanism, using a polymer network with "switching" segments that soften at a specific temperature, allowing the material to be temporarily reshaped, and "netpoint" segments that hold the memory of the original shape and pull the material back once heated.

One detail that trips people up about shape-memory alloys: most exhibit what's called a one-way shape-memory effect, meaning the material remembers exactly one "trained" shape, springs back to it when heated, and needs to be manually deformed again by an outside force to change shape a second time. It doesn't cycle back and forth on its own. A smaller subset of alloys can be processed to show a two-way effect, cycling between two distinct remembered shapes as they're heated and cooled repeatedly, but that requires a more involved training process during manufacturing: repeatedly cycling the material through controlled heating and deformation until it "learns" both shapes. Most wearable applications stick with the simpler one-way version, which is enough for a single functional transition (loosening a cuff, releasing tension, adjusting shape once) rather than continuous, repeated shape-cycling.

In apparel, these properties translate into useful applications rather than gimmicks. Compression garments used in medical and athletic recovery contexts can use shape-memory polymers to apply variable pressure that adjusts with body temperature. Orthopedic and prosthetic-adjacent wearables use shape-memory alloys for components that need to shift shape in response to the wearer's body heat, with no motor or battery involved. Some experimental adaptive-fit footwear and outerwear research has explored garments that physically loosen or tighten in response to temperature or activity level, entirely through the material's own thermal response rather than any embedded electronics.

PCMs and shape-memory materials run on entirely different physics, but they share something: both push apparel from a passive, static material toward one with a built-in response function, clothing that reacts to a condition without a battery, sensor, or microcontroller telling it to. That's a different design philosophy from the "smart clothing" most people picture when they hear the phrase, which usually assumes embedded electronics. These materials get "smart" behavior for free, baked directly into their molecular or crystalline structure.

THE CHEMISTRY NOBODY THINKS ABOUT WHEN THEY BUY A T-SHIRT

By Traicy Moren

Somewhere in the life of most colored garments, there's a vat of water that ends up more colorful, and more toxic, than the fabric it dyed. Textile dyeing is one of the most water- and chemically-intensive stages of garment manufacturing, and understanding why means understanding a bit of the chemistry happening in that vat.



The majority of synthetic dyes used industrially, especially for cotton and other cellulose fibers, are azo dyes, defined by a nitrogen-nitrogen double bond ($-N=N-$) linking two aromatic ring structures. This bond is what gives azo dyes their color: it creates an extended system of alternating single and double bonds called a conjugated system, which absorbs specific wavelengths of visible light and reflects the rest, producing the color you see. Azo dyes are popular because they're cheap to synthesize, produce a huge range of vivid colors, and bond well to fiber. The chemistry problem shows up later. Under certain conditions, including reductive conditions that can occur in wastewater treatment, some azo dyes break down at that $N=N$ bond and release aromatic amine compounds, a number of which are recognized carcinogens. That's why the European Union restricts a specific list of azo dyes known to cleave into these problematic amines, under REACH chemical safety regulation.

Dyeing certain fibers, particularly wool, also often requires a mordant, a chemical that helps bind dye molecules more permanently to the fiber. Historically, and still fairly commonly, that mordant is a chromium compound, because chromium ions form strong coordination bonds with both the dye molecule and the protein structure of wool fiber. The issue is that hexavalent chromium, a byproduct that can form during this process, is both toxic and a recognized carcinogen, and managing it responsibly requires

wastewater treatment infrastructure that not every dyehouse has.

There is an older alternative that predates all of this synthetic chemistry: natural dyes extracted directly from plants, insects, or minerals, the way textiles were colored for most of human history before the first synthetic dye, mauveine, was accidentally discovered in 1856. Natural dyes avoid the aromatic amine problem entirely, but they come with their own limitations. Many still require a mordant to bind properly to fiber, the color yield per unit of raw plant material is often low enough to make large-scale production land- and water-intensive in its own right, and batch-to-batch color consistency is far harder to control than with a synthesized molecule manufactured to a precise chemical specification. That trade-off is part of why synthetic dye displaced natural dye so completely and so quickly once it became available.

Scale is what turns this from a chemistry footnote into a global environmental issue. Textile dyeing is considered one of the largest industrial water consumers and polluters worldwide, and a significant share of that impact comes specifically from the dyeing and finishing stage rather than fiber growth or garment construction. Conventional dyeing is also water-intensive by its basic mechanism: dye typically needs to be dissolved and applied in an aqueous bath, and a considerable volume of water is needed per unit of fabric, most of which b

ecomes wastewater carrying unfixed dye, salts, and processing chemicals.

That dependence on water is exactly what a newer approach tries to eliminate. Waterless dyeing technology, commercialized by companies including DyeCoo, uses supercritical carbon dioxide instead of water as the medium that carries dye into fiber. Above a specific temperature and pressure, its critical point, CO_2 stops behaving like a distinct gas or liquid and becomes a supercritical fluid with properties of both: it diffuses like a gas but can dissolve substances like a liquid, making it an effective solvent for dye without needing water at all. The CO_2 itself can then be captured, depressurized, and reused in a closed loop, and because there's no water bath, there's almost no dye-laden wastewater to treat. Nike and IKEA have both invested in and adopted this technology for portions of their supply chains.

None of these fixes are simple industry-wide swaps. Supercritical CO_2 systems require specialized, expensive equipment and currently work best on synthetic fibers like polyester, not the full range of natural fibers the industry relies on. Regulating azo dye chemistry doesn't eliminate demand for cheap, vivid color; it just shifts formulation toward compounds that don't cleave into carcinogens, and manufacturers in less-regulated markets don't always have the incentive or infrastructure to adopt those compounds. The chemistry of a dyed shirt is an unfinished problem, sitting behind color choices most people never think twice about.



Popularized by designer Christian Dior in 1947, spotted cat prints remain the most enduring animal pattern in modern clothing.

WHAT SHARKSKIN AND LOTUS LEAVES TAUGHT TEXTILE ENGINEERS

By Ferozh Shakouri



For a few years in the 2000s, competitive swimmers were quietly outsourcing their aerodynamics research to sharks. Fastskin suits were textured with tiny overlapping ridges modeled on the dermal denticles that cover a shark's skin, engineered to manage turbulent water flow the same way those denticles do for the animal. The ridges don't smooth the surface the way you'd expect speed to demand. Instead they disrupt turbulence at a microscopic level, breaking up the vortices that would otherwise create drag. It worked well enough that swimming's governing body eventually restricted full-body suits like these over fairness concerns: a strange kind of scientific validation, something banned because it worked too well.

The manufacturing side is less glamorous than it sounds. You can't 3D-print a shark's actual dermal denticle structure onto swimwear at any practical cost. The real approach was closer to embossing a fine, ridged pattern into a thin polymer film, then laminating that film onto the fabric surface. That's a pattern across biomimetic textile manufacturing generally: the biological original is rarely copied atom for atom. It's approximated with whatever fabrication process (embossing, laser etching, coating) can hit close enough to the original geometry to reproduce the effect cheaply and repeatably, across meters of fabric rather than square centimeters of leaf or wing.

That's biomimicry in short. Engineers look at a structure evolution has already optimized over millions of years, figure out the physics behind why it works, and try to rebuild it at a scale and material that's actually manufacturable. Textiles are an unusually rich place to do this, because so much of what fabric needs to do (repel water, resist tearing, regulate temperature, look good doing it) has already been solved somewhere in the biological world.

Take the lotus leaf. Touch one after rain and water doesn't sit on it. It beads into near-perfect spheres and rolls straight off, carrying dirt particles with it. Botanist Wilhelm Barthlott identified why in the 1990s: the leaf's surface is covered in microscopic bumps, coated in a layer of waxy crystals, that dramatically reduce the actual contact area between a water droplet and the leaf. The relevant measurement here is contact angle, how steep an angle a droplet forms against a surface, and the lotus leaf routinely exceeds 150 degrees, the threshold materials scientists use to define superhydrophobicity. Textile engineers have spent two decades trying to replicate that same micro- and nano-scale bumpiness on fabric fibers, chasing self-cleaning clothing that never quite needs washing. The tricky part isn't the chemistry, it's durability: those delicate surface structures wear away with normal abrasion, so a jacket that resists water perfectly on day one might lose most of the effect within a few months of use.

Then there's the example that gets cited constantly, and for good reason: spider silk. Dragline silk, the thread a spider uses for a web's structural spokes, has tensile strength comparable to some steels by weight, but it also stretches significantly before breaking, giving it a toughness that exceeds Kevlar. That combination comes from its molecular architecture: stiff, crystalline beta-sheet regions embedded within a stretchier, amorphous protein matrix, so the material gets a rigid material's strength and a rubbery material's give at the same time. Spiders can't be farmed for silk at industrial scale (they're territorial and will eat each other), so companies have turned to synthetic biology instead, engineering bacteria or yeast to produce spider silk proteins and spinning those proteins into fiber. Japan's Spiber Inc. has been doing this for over a decade, and its lab-grown silk has already shown up in outdoor gear collaborations.

Biomimicry isn't just aesthetic borrowing. Natural selection runs an extremely long, well-resourced optimization process, and it's often cheaper to reverse-engineer an answer that already exists than to derive one from first principles. Biology optimizes for an organism's survival, not for manufacturability, so a lot of biomimetic textile research comes down to translation: taking a structure that works beautifully at the scale of a single leaf or a single spider, and figuring out how to fabricate it by the kilometer, cheaply enough that someone will actually wear it.