

Technology Should Adapt to Us—not the Other Way Around

The moral case for bio-personalized, bioadaptive technology

By Hannes Bend · August 4, 2026 · 14 min read



Concept visualization: living systems continuously respond to changing conditions. A bio-personalized interface applies that principle to everyday technology while keeping the person in control.

What is bio-personalized technology?

Bio-personalized technology uses consented physiological, behavioral, and experiential context to adapt an interface or environment to an individual. Unlike conventional personalization, which primarily learns from clicks, purchases, and past preferences, bio-personalization can respond to present-state signals such as breathing, pulse trends, posture, visual behavior, or self-report. Its purpose should be to support agency, accessibility, wellbeing, and performance—not to increase manipulation or time on screen.

Everything alive is adapting.

Our pupils adjust to light. Our breathing changes as we move, rest, concentrate, become afraid, or feel safe. Trees grow toward light. Rivers reshape themselves around resistance. Nervous systems respond before we have found words for what we feel.

Yet one of the environments we spend the most time inside—the screen—still behaves as though every human nervous system were the same.

The same brightness.

The same typography.

The same notification rhythm.

The same colors.

The same information density.

The same endless scroll.

We call technology personalized because it remembers what we clicked, bought, watched, or searched for.

But that is not deeply personal.

It usually knows our behavior better than it knows our wellbeing.

I believe the next meaningful stage of technology is not simply more personalization. It is **bio-personalization**: technology that can—with our permission—understand something about our present physiological state and adapt its visual, auditory, and interactive environment around us. This is the practical face of what I call [Somatic AI](#).

to ourselves.

I did not arrive here through technology alone

My path into adaptive interfaces did not begin with an AI trend.

It began with my body.

I was born with mild cerebral palsy and nearly died at birth. For much of my early life, I struggled with physical pain, anxiety, periods of depression, and a feeling that life was something heavy I had to push through.

I first found a different reality through yoga and meditation. Later, breathing awareness taught me that a shift in physiology could become a shift in perspective.

One deeper inhale.

One letting-go exhale.

One moment in which the world was no longer exactly the same.

That experience became a question that has followed me for years:

If a breath can change how I experience reality, what else could help people return to a more regulated, aware, and capable state?

I was an **artist** before I was an AI founder.

Art taught me that perception is not passive. Color, scale, movement, material, light, and context do something to us. Two people can stand before the same work and have completely different physical and emotional responses.

There is no universally calming image.

There is no universally inspiring color.

There is no single interface that is optimal for everybody.

That is not a problem to be designed away. It is an essential part of being human.

Metaverses: asking what images do to us

In 2014, I became a visiting scholar at the University of Oregon's Alemán Quantum and Nanoscale Physics Lab and an artist-in-residence connected to the Institute of Neuroscience.

My [Metaverses](#) / [Seecular](#) research asked a question that felt surprisingly underexplored:

What happens in the brain and body when we look at different images and image flows?

We recorded EEG activity while participants viewed and rated tens of thousands of images. We explored relationships between subjective preference, visual aesthetics, attention, emotion, and patterns of brain activity associated with calmer or meditative states.

The purpose was not to find the one image that would calm everybody.

We did not find one.

The more important insight was that people responded differently.

That difference is the foundation of personalization.

Our later [Mindful Technologies](#) work connected these questions to virtual reality, breathing, pulse, biofeedback, sound, and immersive environments—work that continued through [SEAing Breath](#) and the mYndful prototypes now documented on the [research page](#). The virtual environment could begin to respond to the participant rather than simply play a predetermined sequence.

The person was no longer only a viewer.

The person's body became part of the medium.

A retinal implant changed how I thought about interfaces

While I was at the Alemán Lab, another interdisciplinary University of Oregon team led by physicist Richard Taylor was developing a bio-inspired retinal implant.

The research used fractal geometry—patterns found throughout nature, including in neurons—to explore how artificial electrodes might interface more effectively at or near the retina, interacting with surviving retinal neurons.

The project received a major [W. M. Keck Foundation grant](#). It was led by Richard Taylor's team, not by me; my own work at the university sat in the neighbouring [Alemán Lab](#).

For me, the idea opened a much larger question.

If artificial systems could one day restore aspects of vision by delivering signals directly through retinal interfaces, then visual design was no longer only something placed on a screen across the room.

It could become much more intimate.

It could become part of what a person is physically able to perceive.

And if a visual signal can be introduced so close to the nervous system, the quality and intention of that signal become even more important.

What should we place there?

Who decides?

Should the output be identical for everybody?

Could it adapt to a person's abilities, sensory processing, state, or needs?

Could a font respond to changing eyesight?

Could contrast adapt without an older user searching through layers of settings?

Could an interface become more legible before the user has to admit they are struggling?

Could visual or auditory output help a person regulate rather than overwhelm their nervous system?

These questions eventually entered my [adaptive-interface patents](#).

continuously personalized output: color, font, brightness, contrast, saturation, graphical interfaces, audio volume, pitch, and other elements [2, 3].

The technology is broad.

The moral purpose must be clearer.

How this developed

2011

Yoga and meditation begin changing Hannes Bend's relationship with his body and mental wellbeing.

2014

Breathing awareness; visiting scholar and artist-in-residence work at the University of Oregon.

2014–2016

Metaverses / Secular research into visual stimuli, EEG, meditation and perception.

2016

Mindful Technologies presented through the AAAI Spring Symposium.

2016

University of Oregon retinal-implant work receives the W. M. Keck Foundation grant. The retinal-implant project was led by Richard Taylor's interdisciplinary University of Oregon team and included Benjamin Alemán.

2016

The widely circulated VR-headset audience photograph becomes a cultural symbol of both the promise and the isolation risk of immersive technology.

2017

Adaptive-interface patent priority filings.

2018

MIT-area and laptop-based adaptive-interface prototypes.

2019 AND 2023

U.S. patents granted.

CURRENT

Bio-personalized interfaces connecting physiology, AI and everyday technology.

A personal chronology. The retinal-implant project was led by Richard Taylor's interdisciplinary University of Oregon team and included Benjamin Alemán. A longer version of this chronology is on the [progression page](#).

VR showed me the possibility—and the limitation

I still believe virtual reality can be powerful.

Our early heart-rate and breathing biofeedback experiences made internal biological processes perceptible. Breath could become sound or movement. Pulse could become part of a visual world. VR could teach interoception in ways that a numerical health dashboard could not.

But VR also requires another device, another application, another deliberate experience, and usually more time inside technology.

I remember the widely circulated 2016 photograph of Mark Zuckerberg walking through a room filled with people wearing VR headsets. Nobody around him appeared aware that he was walking past.

The image was more complicated than the dystopian interpretations it received. The participants were sharing a VR experience, and virtual environments can create connection, education, empathy, and accessibility.

But the photograph still stayed with me.



Immersive technology can create connection and access, but it can also narrow awareness of the physical people and environment around us. The illustration references the debate triggered by a widely circulated 2016 Mobile World Congress photograph. The reference is to the photograph taken at Mobile World Congress, Barcelona, February 2016, widely reproduced in technology press coverage at the time.

It represented a possible future in which the digital environment replaces the entire visual field while the physical people and world around us become less perceptible.

I did not want my work to depend only on putting more screens in front of people's eyes.

When I returned from research environments to New York, the bigger problem was obvious.

People were already spending enormous amounts of time looking at phones and computers.

The intervention had to begin where people already were.

Inside the email.

Inside the chat.

Inside the browser.

Inside the work interface.

Inside the pixels people were already required to look at.

The interface is already influencing us

There is no neutral interface.

Every interface directs attention.

Every default makes one behavior easier than another.

Every notification asks to become more important than whatever we were doing one second before.

Every feed determines what appears next.

Every design system creates a rhythm for the nervous system interacting with it.

Today, many platforms are optimized around engagement, retention, reaction, and conversion.

Unpredictable social feedback—new likes, messages, alerts, and content—can reinforce repeated checking. It is too simplistic to say that scrolling merely “releases dopamine.” Dopamine is involved in learning, motivation, anticipation, reward prediction, and reinforcement, not simply pleasure.

But the larger point remains:

Digital systems have become extremely sophisticated at learning what keeps us returning.

They are much less sophisticated at understanding when returning is no longer good for us.

Research increasingly associates problematic screen use and doomscrolling with anxiety, stress, poorer mood, distraction, and reduced wellbeing [11, 12]. Notifications can increase frustration and task errors [16]. Experimental work also suggests that reducing smartphone use can improve stress, sleep, depressive symptoms, and wellbeing [10]. These are associations and short-term experimental findings, not proof that any particular screen causes any particular disorder.

This does not mean every screen is harmful.

A message can reduce loneliness.

A video call can connect a family.

A digital community can help someone who feels alone in their physical environment.

Cambridge Analytica changed the moral stakes

The Cambridge Analytica scandal made this question impossible for me to ignore.

Personal data gathered through Facebook was used to build profiles and target political messages toward groups of voters. Regulators later documented deceptive data collection and political microtargeting practices [8, 9].

We should be careful about exaggerated claims. Cambridge Analytica did not prove that one image could deterministically control a voter or that every targeted person changed their vote. What the record establishes is psychographic ambition, deceptive collection, and susceptibility to manipulation—not election causation.

But it showed the direction clearly.

Personal data could be used not simply to understand people, but to find their susceptibilities and shape what they saw.



The Cambridge Analytica scandal demonstrated how personal data could be used for political profiling and microtargeting. It did not establish that a single message could deterministically control a voter. Original editorial illustration. Primary sources: [FTC case record](#) and [ICO investigation report](#).

physiological information?

When does someone become more receptive?

When are they agitated?

When are they tired?

What visual intensity keeps them engaged?

Which message produces the strongest biological response?

A bioadaptive interface could support regulation and wellbeing.

The same architecture could also become a more intimate instrument of persuasion.

That is why ethics cannot be added after the technology has scaled.

It must exist inside the design.

Two enormous systems that still barely speak

We now have two enormous technological systems developing beside each other.

On one side, we have interfaces.

Phones, chat applications, productivity tools, social media, video platforms, browsers, games, vehicles, and augmented environments provide increasingly personalized content.

But the output itself often remains surprisingly static.

People still have to open settings manually to change font size, contrast, background color, notification behavior, or display intensity.

This is tedious for everybody.

For an older person, someone with changing eyesight, a migraine, sensory sensitivity, fatigue, or reduced dexterity, it can become a real accessibility barrier.

On the other side, we have biodata.

Wearables and sensors continuously collect steps, sleep, pulse, breathing-related data, activity, recovery, temperature trends, and other measurements.

That data is valuable. But much of it remains inside dashboards, and a considerable share of devices are eventually abandoned [23].

It is measured.

It is visualized.


Sometimes it is scored.

But it rarely changes the environment that helped create the state in the first place.

A watch may tell us that our stress indicators are elevated while the screen continues flashing, interrupting, crowding, and demanding attention in exactly the same way.

The two systems barely speak.

Bio-personalization connects them.



Two systems currently develop beside one another: expanding streams of biodata and largely static interface output. Bio-personalization connects them through a consented, explainable, and reversible adaptive loop.

Start with something ordinary: a chat

Consider how much of life now happens through text.

Work conversations. Relationships. Family messages. Political arguments. Medical communication. News. Negotiations. Misunderstandings.

A bio-personalized chat interface would not need to manipulate the content of the conversation.

It could adapt the conditions around it.

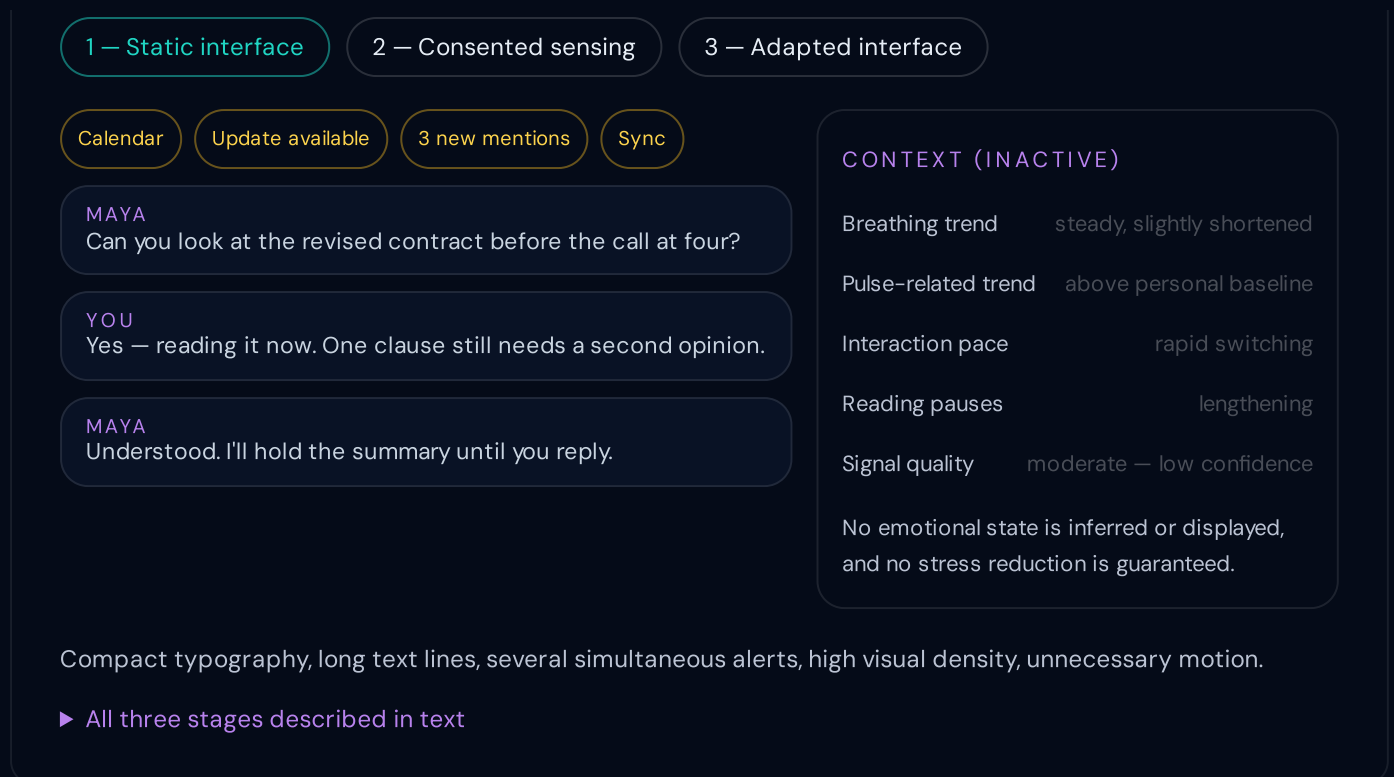
With explicit consent, it might learn that a particular person reads more accurately with slightly larger text and more line spacing.

- It might reduce animation when interaction becomes more demanding.
- It could batch nonessential notifications instead of presenting each one as an emergency.
- It could adjust brightness or saturation modestly.
- It could offer an optional pause after rapid switching between emotionally intense conversations.
- It could make the interface easier to use as eyesight changes.
- It could explain every adaptation and allow the user to reverse it immediately.

This sounds small.

But billions of small digital interactions form a large part of modern life [24].

A slight improvement repeated across thousands of hours may matter more than an extraordinary ten-minute experience requiring a special headset.



Concept visualization: a chat interface adapts presentation and interruption timing—not the user's words—through consented physiological and interaction context.

What the research supports

Research already gives us several pieces of this future.

Different fonts improve reading performance for different individuals. There is no universal best digital typeface [13]. I explored this in detail in [how bioadaptive interfaces can personalize color, typography and interaction](#).

Brightness, saturation, and hue can influence reported emotional arousal, but context and individual response matter [14].

Notification interruptions can increase stress, annoyance, anxiety, and errors [16]. Batching notifications can reduce stress without requiring a complete disconnection from other people [17].

Physiologically adaptive visual environments have been shown experimentally to affect comfort, workload, and working-memory performance [15].

Nature exposure is associated with stress recovery and attention restoration [18]. Virtual nature can also be beneficial, especially where access to physical nature is limited [19].

But science does not yet justify every claim I would like to test.

We do not know that a calmer interface will automatically make someone kinder.

We do not know that a lower heart rate will always produce a better decision.

Acute stress can impair cognitive flexibility in some conditions [21], while other experiments find that moderate stress can improve certain forms of long-term reward learning [22].

Human beings are not simple state machines.

That is why a bioadaptive system should not impose one universal assumption.

It should learn modestly.

Adapt carefully.

Represent uncertainty.

And remain answerable to the person using it.

AI, affective computing and bioadaptive interfaces, and the sensing layer itself is described on the [rPPG API page](#).

My hypothesis: calmer technology can widen our horizon

What follows is my hypothesis and proposed research direction, not an established finding. I believe a more regulated digital environment can create better conditions for reflection.

Not guarantee it.

Create conditions for it.

When we are being interrupted, provoked, hurried, and pulled between competing signals, our decision horizon can become smaller.

We react.

We defend.

We click.

We move to the next thing.

A calmer interface may give us another fraction of a second.

Enough to notice how we feel.

Enough to reread.

Enough not to send the message.

Enough to remember the person on the other side.

Enough to consider whether the decision serves only the next minute or also the next year.

This is a hypothesis worthy of serious research:

Can bio-personalized interfaces support not only individual performance and wellbeing, but more reflective, long-term, and prosocial decision-making?

The answer may not always be yes.

But the question is too important not to test.

The moral architecture of bio-personalization

A bio-personalized interface should not secretly decide what emotional state a person ought to have.

Calm is not always the correct state.

Sometimes urgency is appropriate.

Anger can contain valuable information.

Grief should not be optimized away.

Excitement should not be flattened into an average biometric score.

The purpose is not to normalize people.

The purpose is to support agency.

That requires clear principles:

1. CONSENT

Physiological sensing must be knowingly activated.

2. PURPOSE LIMITATION

Data collected for wellbeing must not quietly become advertising, employment evaluation, insurance pricing or political targeting data.

3. MINIMAL DATA

Use the smallest amount of information necessary.

4. LOCAL PROCESSING

Where possible, process raw camera or sensor information on the user's device.

5. UNCERTAINTY

Do not turn an ambiguous signal into a confident emotional diagnosis.

6. EXPLANATION

The person should be able to understand why an interface changed.

7. REVERSIBILITY

Every adaptation should be easy to pause, reject or restore.

8. NO ENGAGEMENT TRAP

A successful adaptation may help the person complete a task and leave the screen sooner.

should not be measured only through increased time on screen.

Technology should help us return to the world

The World Health Organization now estimates that loneliness affects around one in six people globally [20].

We have never had more tools for connection.

And yet connection is not the same as contact.

A thousand digital reactions do not necessarily replace one person being fully present.

I do not want bio-personalization to create a screen so perfectly tailored that nobody wishes to leave it.

That would be another form of captivity.

I want technology that helps people get what they need with less friction and less physiological cost.

Technology that helps us complete the task.

Understand the message.

Regulate before responding.

Notice our breath.

Then look away.

Go outside.

Be with people.

See the changing light.

Feel weather.

Walk through a natural environment whose complexity no screen fully reproduces.

Virtual nature can be restorative, particularly where physical nature is inaccessible. Biofeedback VR can be meaningful. Digital environments can offer access where physical nature is unavailable. Real nature and physical community remain broader multisensory experiences [19].

But the highest goal is not to replace life with a better simulation

The future I want to build

I imagine interfaces that become more legible as our eyesight changes.

Work environments that understand when another interruption would do more harm than good.

Audio that adapts to hearing, context, and physiological response.

Educational environments that adjust pace without shaming the student.

Chat applications that make understanding easier rather than conflict more addictive.

Digital systems that recognize when they are no longer helping.

Technology that learns how to support us without claiming to know us completely.

The screen becomes less like a slot machine and more like a responsive environment.

The interface is not trying to extract another minute.

It is trying to return something.

Attention.

Clarity.

Time.

Agency.

Maybe even joy.

This is why I developed adaptive-interface technologies, and why they are now being applied through the [ventures](#) I work on.

Not because people need more technology.

Because technology needs to become more human.

The highest achievement of bio-personalization will not be keeping us on a screen longer.

It will be helping us need the screen less.

Technology should meet us where we are, help us return to ourselves, and then know when to

“[stop the car](#)”

Related reading: [From immersive art to remote photoplethysmography](#) traces the same arc through the research record.

Research and sources

1. Hannes Bend et al., "Immersive Mindful Technologies" (AAAI Spring Symposium, 2016) — PDF
2. Hannes Bendfeldt, "Adaptive Interface for Screen-Based Interactions" — US10423893B2
3. Hannes Bendfeldt, "Adaptive Interface for Screen-Based Interactions" — US11561806B2
4. University of Oregon, "UO idea: bio-inspired implant wins \$900,000 grant" (W. M. Keck Foundation)
5. Alemán Lab, University of Oregon — group members
6. Hannes Bend, "Your UI Is Here For You"
7. Hannes Bend — writing on Medium
8. U.S. Federal Trade Commission — In the Matter of Cambridge Analytica, LLC
9. UK Information Commissioner's Office — Investigation into the use of data analytics in political campaigns (2018)
10. Experimental smartphone-use reduction intervention and mental-health outcomes
11. Doomscrolling, news consumption and wellbeing
12. Problematic screen use and psychological outcomes — review evidence
13. Shaun Wallace et al., "Towards Individuated Reading Experiences: Different Fonts Increase Reading Speed for Different Individuals"
14. Wilms and Oberfeld, "Color and emotion: effects of hue, saturation, and brightness"
15. Francesco Chiossi et al., "Adapting Visual Complexity Based on Electrodermal Activity Improves Working Memory Performance in Virtual Reality"
16. Notification interruptions, stress and task performance
17. Batched notification delivery and wellbeing outcomes
18. Nature exposure and human health — review evidence
19. Comparing actual and digital nature for restorative outcomes
20. World Health Organization — Commission on Social Connection report
21. Acute stress and cognitive flexibility
22. Stress and long-term reward learning
23. Wearable device abandonment — analysis
24. DataReportal — Digital 2026 Global Overview Report

Image credits: all illustrations on this page were created for this article. No historical news photograph or political advertisement was reproduced; source links are provided instead.