ARE IN AI ™

San Diego's First Responsible & Human-Centered AI Forum

Human-Centered AI Pulse

A snapshot of the inaugural San Diego Responsible & Human-Centered AI Forum: the themes that emerged and how we can intentionally shape what comes next.

WE ARE IN AI™ San Diego Human-Centered AI Pulse 2026
Published October 2026 |CC BY-NC-ND 4.0

www.weareinai.org



CONTENTS

TABLE OF CONTENTS

1	The Inflection Point	02
2	AI Landscape	03
3	The Opening Panel	04
4	Closing Reflections	07
5	Key Themes	08
6	Closing Plenary	10
7	Calls to Action	11
8	Voices from the Forum	13
9	The Convergence	14
10	Topics Considered	15
11	Methodology	16
12	The SPHERES	17
13	An Invitation	18
14	About WE ARE IN AI™	19
15	Acknowledgments	20

CONTEXT

THE INFLECTION POINT

AI Is Not Only a Technology Question, It Is a Human One.

We are living in an AI-native world. Not approaching one. Capital is flowing into data centers at a scale that reshapes regional power grids. AI sits in decisions that used to require a person's judgment, and in the ordinary machinery of everyday work. Very little is untouched. That is the ground this conversation starts on. ***The open question is not whether AI changes things. It is who the change is built for, and whether we will shape that transformation intentionally.***

Across industries, organizations are moving to deploy AI faster than most are actually ready for. Speed is not the same as strategy. When speed becomes the only metric that matters, something quietly gets optimized out of the architecture: the human being this was all supposed to serve.

AI is often described as a foundational technology, the kind that other systems and decisions get built on top of, similar to how electricity or the internet became infrastructure for everything that followed. A design choice made early doesn't stay contained: it becomes part of what everything built afterward inherits. Human-centered AI means asking who a system serves before asking how fast it can launch, a question that gets answered at the design stage whether anyone raises it or not. This isn't new. Decades of human-centered design have already shown that systems built with people outperform systems built without them. Human-centered AI applies that same principle to the highest-stakes technology yet.

> The choices we make today about AI will echo through generations. The future of AI cannot be inherited, it has to be designed, and it has to be designed intentionally, with people at the center.

WE ARE IN AI™ is a mission-driven organization and global ecosystem for Human-Centered Innovation and Responsible AI. As part of our mission, we convened San Diego's first Responsible & Human-Centered AI Forum the region's inaugural forum of its kind, on one premise: building human-centered AI cannot belong to one industry or discipline. It has to start with all of us, business, technology, academia, and community, working toward a shared vision through collaboration.

The organizations and regions that shape AI's future won't be the ones with the most technology. They'll be the ones asking who it serves. Most AI conversations start with the technology. Ours started with people. The day was facilitated and structured using the WIN AI™ methodology and its SPHERES dimensions for human-centered AI, applied across the Business, Ecosystem, and Community tracks.

The conversation about human-centered AI is too important to end when the forum did. We wanted it to inspire something larger: ongoing thinking, conversation, and collaboration to actively design the future, not just react to it. That is why this document exists. It is released as the first San Diego Human-Centered AI Pulse 2026, a living document meant to keep the conversation going.

San Diego's Human-Centered AI conversation started here.

CONTEXT

AI LANDSCAPE

What's Already Happening in San Diego

The City of San Diego started early. It adopted a formal AI policy and established AI governance processes in 2024, at a point when few municipal playbooks existed. CIO Jonathan Behnke led the rollout of AI productivity tools across city operations beginning with a structured pilot and expanding to all 13,000 city employees. Within two days of opening AI training to all employees, 1,800 had already signed up.

8th
Largest city in America

1.4M
Residents

13,000
Employees

Why Community Engagement Matters Now

The speed of AI deployment makes community engagement essential. Decisions about which processes to automate, which data to use, and which communities to serve first are being made today, faster than oversight can keep up with. Those decisions aren't made by government alone. They're made by the companies building the models, the organizations choosing what infrastructure to invest in, the employers deciding on roles and automation, and the vendors whose default settings quietly become everyone's standard practice.

Most of that happens in rooms residents and workers never see, long before any regulation catches up to it. Who can participate is uneven: Pew's American Trends Panel analysis found that higher-income and more-educated adults recognized AI operating in ordinary settings more often than others, a divide in visibility rather than device access. Who bears the cost is uneven too: Brookings found that of 37.1 million US workers in the top quartile of AI exposure, 6.1 million have both high exposure and low adaptive capacity, concentrated in clerical and administrative roles, and about 86% of that group are women. And who captures the gain remains an open question: the International AI Safety Report notes that general-purpose AI could widen income and wealth inequality by shifting earnings from labor toward capital owners, including shareholders of the firms developing and deploying it. Community engagement means putting local voices in the room early, while there's still a chance to shape the outcome.

WHERE THE DAY BEGAN

THE OPENING PANEL

The forum opened with a panel discussion featuring voices from government, HR and workforce, AI ethics and design, and labor, a cross-sector conversation about what responsible AI adoption actually looks like from inside the institutions deploying it. The panel moved across government policy, workplace disruption, regulation, and personal reflection, but kept circling back to the same tension: efficiency gains that are genuine, sitting next to unresolved questions about who those gains actually serve. Guiding questions included:

- What does human-centered and responsible AI actually mean?
- What's the current state of AI in your industry?
- What's the most important question about AI that organizations and communities aren't asking?
- Where should AI go in the next few years, if we get it right?
- What should people leave today thinking or doing differently about AI?

> "The costs, the displacement, the consequences of AI are landing on people who had no say in the systems affecting them." - *From the panel discussion*

Three Layers of Meaning

The panel started by setting the stage and defining key concepts. Di Le laid out three ways to think about this: ***Ethical AI*** is our moral compass, the philosophical layer, society's values made explicit. ***Responsible AI*** is the practical layer, turning that philosophy into frameworks, guidelines, and governance. ***Human-Centered AI*** sits on top of both: it's the human-impact lens, asking not just "does this comply?" but "what do humans take pride in, and what do we actually want automated?" This distinction matters: compliance and humanity are not the same thing.

Responsible AI Governance in San Diego

The City's AI story started in 2023, right after generative AI's public debut in late 2022, with a cross-departmental policy team building the City's AI policy from early NIST and outside industry research, at a time when almost no government playbook existed yet. What began as a proof-of-value pilot in January 2025 scaled citywide within nine months.

The Gov AI Coalition, which San Diego helped found, is a case study in how fast shared standards can scale once someone starts building them: A single meetup of eight CIOs has grown into a network of more than 900 state, local, and federal member agencies and 3,000 practitioners nationally. Members built shared policy templates and "vendor fact sheets," a nutrition-label-style comparison tool for AI vendors. Despite spanning a genuinely broad ideological spectrum, the coalition has never had anyone push back on the basic premise of responsible AI. San Diego aligned its own AI policy with its existing trust ordinance that requires community feedback across all nine council districts before any new citywide deployment.

The Workplace Is Already Transforming

On the HR and workforce front, AI has quietly rewritten pay transparency. Workers can now get a precise figure by role, experience, city, and even a company's funding stage. The upside is real: workers know their worth. The downside is a pay-compression problem now facing HR teams everywhere, as new hires arrive with data-backed demands and existing employees discover the gap.

On the regulatory side, at the time of this writing, state AI employment law is moving from disclosure toward audit, anti-discrimination, and human-review rights, though unevenly and on different timelines depending on the state. Illinois's HB 3773 prohibits AI that has a discriminatory effect and requires notice. Colorado's revised AI Act goes further, requiring an adverse-outcome explanation and an offer of meaningful human review, once it takes effect. California's Civil Rights Council regulations focus on anti-discrimination, not human-in-the-loop decisions. NYC's Local Law 144, a city rather than state measure, mandates bias audits and disclosure. None of them currently requires a human to make every final call, though Colorado comes closest, giving individuals the right to request one.

The Governance Gap and the Democratic Imperative

The core tension named on the labor and regulatory side: externalities, costs borne by people who had no say in the transaction creating them, need to be priced into the cost of doing business. A factory doesn't get to dump its waste into a river without paying for the damage. AI shouldn't get to treat displaced workers and eroded trust as costs no one has to answer for. Regulation is genuinely in flux: multiple executive orders have already been issued, at least one already rescinded, with more open questions than settled answers this early in the technology's life.

The case made on the panel: these tools are too powerful to leave in the hands of a few private companies. Some form of democratic control, with real protection for those with the least power to shape the outcome, is worth building toward, even without a finished blueprint yet. A related gap showed up closer to home: leadership is increasingly mandating AI use before the training or governance to use it responsibly exists. The problem isn't that regulation is behind the technology. It's that organizations are outrunning their own capacity to deploy it well.

> "Make human and philosophical conversations important again. Treat the question of how to live a good life in the age of AI with the same seriousness as any technical one." - *From the panel discussion*

Actions Over Statements

"We should trust what people do, not just what they say" - From the panel discussion. The observation behind it: major AI companies have publicly asked to be regulated while simultaneously building data centers at scale and declining to fully explain how their systems work. One panelist framed the fix as carrot-and-stick: the obligation to explain your own systems comes before you get to ask the public for trust and data.

A Global Conversation

While this room was meeting in San Diego, governments elsewhere were answering the same questions with different tools, and the picture below reflects where things stood as of this writing.

The European Union began enforcing the world's first comprehensive AI regulation in August 2026, just weeks before this forum took place, built around a staggered set of obligations: transparency requirements first, followed by rules for high-risk systems, then for AI embedded in regulated products. The transparency rules reach further than most organizations realize: anyone putting their name on a chatbot has to disclose that people are talking to a machine, and the rules apply outside the EU as well.

Asia isn't waiting on Europe. South Korea's AI Basic Act covers high-impact systems with obligations for risk assessment, transparency, and human oversight, and reaches foreign providers too. China requires visible and embedded labeling of AI-generated content, one of the strictest content-provenance regimes in place anywhere. Japan took the opposite approach, passing AI legislation with no monetary penalties at all, betting that binding rules would put its own developers at a disadvantage.

The United States still has no comprehensive federal AI regulation. What exists instead is a mix of state-level rules taking shape at different speeds, with states and the federal government still working out how much authority belongs where. California has gone a step further than regulation alone: in May 2026, the state launched Engaged California, a platform inviting residents to weigh in directly on how AI is affecting their work and communities, feedback the state says will help shape its policy response.

Underneath the differences, one thing converges. Every serious effort to regulate this has landed on the same minimum: tell people when they are dealing with AI. Disclosure is becoming the floor that everyone agrees on, whatever else they disagree about.

And the question is being asked outside of regulation entirely. Pope Leo XIV's first encyclical, released just over a week before this forum, made its own case for what AI owes the people it is built for. Two conversations, half a world apart, arriving at the same question in the same week with no coordination between them. That is what it looks like when something stops being a niche concern for technologists and becomes everyone's question. The argument for human-centered AI is not being made in one forum. It is being made in all of them at once, in rooms that have never talked to each other.

The Utility Question

An AI CEO has said publicly he wants intelligence treated as a utility, priced like electricity. That idea came up directly on the panel. The room answered fast: then make it public, not privately metered, built to serve everybody, not just whoever can pay. Underneath that exchange sat a harder question: **is there even intelligence here to price?** One panelist called AI, as commonly used, a marketing term for a set of tools, not intelligence in any real sense. If that's right, the entire utility debate is arguing over how to meter something whose name may not describe what it actually is.

> “For now AI can’t create what we haven’t already imagined. It recombines what’s already been made. The imagining still has to happen first, and that part still belongs only to us.” - *From the panel discussion*

The Capability Gap

Every technological disruption eventually settles into standards and regulations, usually years behind the technology itself. That gap is the argument for building protection in now, not after the fact. None of this resolves the panel’s central tension. It sharpens it. Real efficiency, sitting next to real questions about who it serves. The work this room started wasn’t answering that question. It was agreeing to keep asking it together, before the answer gets decided for everyone.

CLOSING REFLECTIONS

When a big new technology shows up, the obvious thing everyone expects to change is the work, the tasks, the tools, the job itself. Something more surprising changed first: catching yourself trusting a machine more than you should, second-guessing your own writing, wondering if the strengths that got you here still count for as much as they used to. That's people quietly taking inventory, working out in real time exactly why keeping humans at the center of this technology matters.

One line looked backward to look forward: the last time a technology this powerful moved this fast, it was nuclear energy, deliberately kept out of private hands, not because it lacked promise, but because the risks demanded a kind of restraint that speed alone couldn't provide. “Move fast and break things” already had its turn. That same restraint is what's missing now, and the stakes reach past this generation.

Today's adults learned how to think, and how to trust their own judgment, from other people. The next generation may learn it primarily from an AI tutor instead. Whether they ever get to experience what human agency actually feels like depends on choices being made now.

Alongside the reflections, the panel's calls to action: support local arts, get involved, and attend your city council meetings. Ask questions out loud even when you don't have the answer. Someone else in the room might have it, or might be sitting with the exact same question, and that's what starts the conversation. Finally, make philosophical conversations cool again, treating the question of how to live a good life in the age of AI with the same seriousness as any technical one.

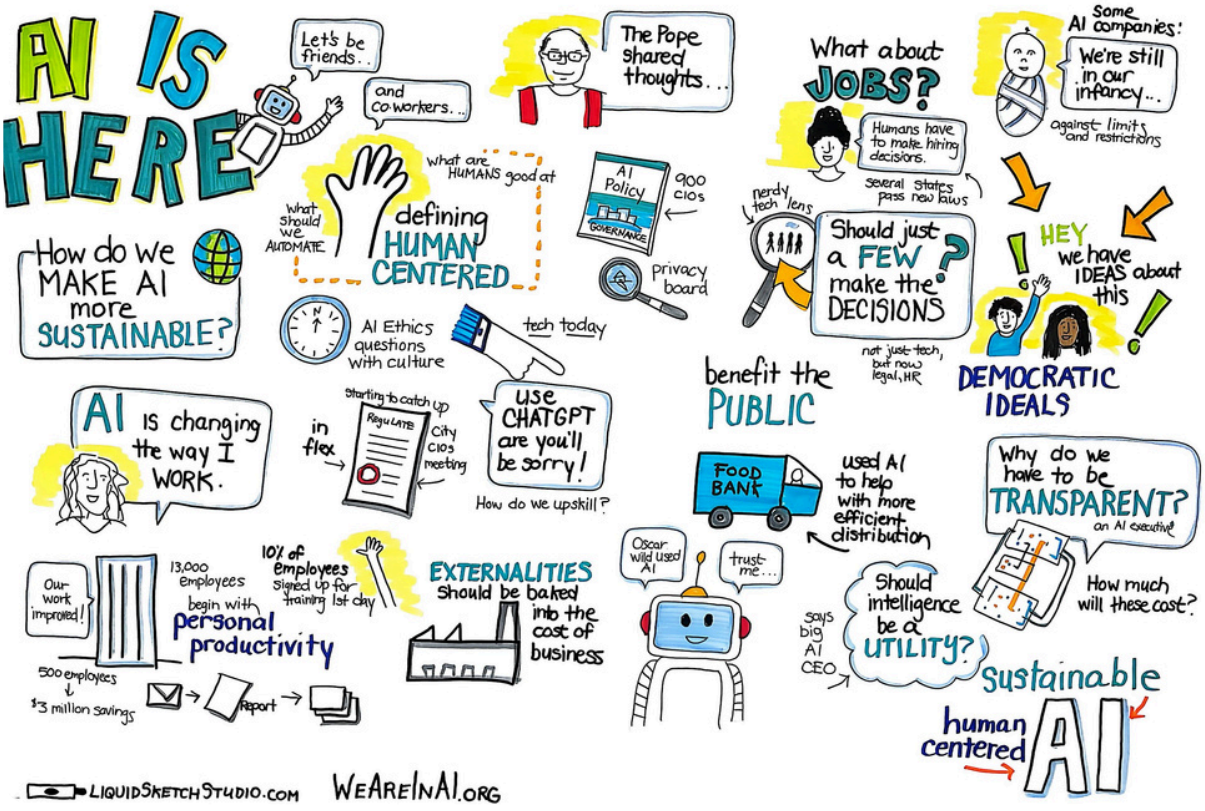

Live graphic recording captured during the opening panel "AI Is Here." The "should intelligence be a utility?" bubble and the Pope reference in the sketch both trace back to real moments in the panel. Illustration by Liquid Sketch Studio (@LiquidSketchStudio) for WE ARE IN AI™. A human-centered AI forum deserves a human record of it. Graphic recording wasn't a visual add-on. It was the same argument the forum itself was making, demonstrated instead of stated: creativity, interpretation, and meaning-making are deeply human capabilities, worth preserving especially in an age of AI.

A CLOSER LOOK

KEY THEMES

After the opening panel, the forum ran three tracks, Business, Ecosystem, and Community. Groups were self-organized that morning, people who hadn't worked together before. Within each track, groups worked through WIN AI™: naming what human-centered and responsible meant for their specific challenge, selecting two SPHERES™ dimensions as their lens, and moving from that understanding to a concrete call to action for the region. Five conversations, one shared process.

BUSINESS

Incentives for AI Adoption

WHAT ARE THE RIGHT INCENTIVES FOR PEOPLE TO USE AI?

> **"Everyone needs a say: designing the right incentives starts with a collaborative conversation, not a top-down mandate."**

This conversation looked at incentive alignment from three angles: the company's interest in revenue, the individual's interest in whether AI helps them do their job, leadership's role in inspiring adoption, and asked whether today's incentives get us there deliberately, or just to next quarter.

CAME IN THINKING
Incentives are mainly a company or individual question with leadership expected to decide how to inspire adoption

NOW BELIEVE
Everyone needs a say: a collaborative discussion across every function and level, with real representation before decisions.

OPEN QUESTION
How do we foster genuinely inclusive conversations and then actually execute on what comes out of them?

CALL TO ACTION
Get educated on AI's consequences. Hold intelligent conversations within the region about what we're building.

BUSINESS

Navigating Disruption

HOW CAN WE MOST EFFECTIVELY NAVIGATE THIS TRANSITION WHEN THE POTENTIAL FOR DISRUPTION IS SO HIGH?

> **"There's far more overlap than expected: everyone needs to be part of shaping AI, not just leadership."**

This group wrestled with the tension between AI's upside and its disruptive cost. A central distinction emerged between "disruptors" and "the disrupted". Human-centered design means involving the disrupted in shaping the solution, not designing around them. "Power of agency" emerged as the core SPHERE Dimension: creativity as the outcome of active participation, not passive absorption.

CAME IN THINKING
Disruption looks different for owners and executives than for average employees; leadership needs to guide people through it.

NOW BELIEVE
The transition requires everyone's active participation, not something leadership manages for others.

OPEN QUESTION
What is actually preventing people's agency? Fear? Lack of education? Systemic barriers? Unequal access?

CALL TO ACTION
Exercise whatever agency you have, whenever it shows up, get involved in decisions about AI in your own processes.

ECOSYSTEM

Maintaining Human Control & Intention

HOW DO WE MAINTAIN HUMAN CONTROL AND INTENTION IN AI?

> **"AI is not one-size-fits-all and it shouldn't be built that way."**

This conversation mapped the full stakeholder ecosystem: parents, mental health professionals, policymakers, founders, educators. From a broad brainstorm spanning education, AI's energy footprint, and representative training data, the group converged on a focused synthesis: the challenge is ultimately an education problem, solved by "baking in" human-centered values from the start.

CAME IN THINKING
Critical thinking, regulation, discernment, and keeping people central are all important.

NOW BELIEVE
AI should never be deployed as a one-size-fits-all solution. Different people, different needs, different guardrails.

OPEN QUESTION
How will people in power be incentivized to account for diverse needs? How will AI literacy reach the people who need it most?

CALL TO ACTION
Always include humans: create forums, community engagement, education and exercise your voting rights. Build guardrails tied to data and policy.

COMMUNITY

Preserving Humanness

HOW DO WE PRESERVE OUR HUMANNESS AND HUMANITY IN GENERAL AS AI ADVANCES?

> **"We came in thinking this was an AI problem. We left knowing it's a human one."**

The group reframed its opening concern: the question isn't how to protect humanity from AI, it's how to exercise the distinctly human capacities AI cannot replicate: critical thinking, creativity, making. The conversation also surfaced that uncertainty, growth, and revision are freedoms worth protecting as AI optimizes toward confident, stable outputs.

CAME IN THINKING
This is a problem specific to AI something the technology is doing to us.

NOW BELIEVE
This is a human problem, our responsibility, not AI's. It starts with self-reflection and ripples outward in what we choose to include, build, or leave out.

OPEN QUESTION
What makes us human and how do we preserve it as AI becomes capable of emulating more of what we thought was distinctly ours?

CALL TO ACTION
Start with yourself. Carry it outward community and business. Treat civic duty as part of the answer.

COMMUNITY

Pacing AI Adoption to Community Needs

HOW DO WE KEEP INNOVATION PACED TO WHAT COMMUNITIES ACTUALLY NEED?

> **"How much is too much and who gets to decide?"**

This group named the tension running through the whole day: AI is growing and being adopted rapidly without being fully understood: not by individuals, organizations, or governments. Their call to action carried the most urgency of any group: act now, deliberately, so that by 2036 San Diego is not caught off guard.

CAME IN THINKING
AI is growing and being adopted rapidly without being fully understood.

NOW BELIEVE
The belief hasn't changed: AI is still moving fast. That's exactly what raises the real questions: where is the line, and who draws it?

OPEN QUESTION
How much AI adoption is too much, before it outpaces our ability to understand, govern, and reverse course?

CALL TO ACTION
Act now, deliberately. Everyone has a vote and a voice. Civic responsibility is AI governance.

REFERENCE

CLOSING PLENARY

Several groups, working independently, landed on the same idea: this is a human problem, not a tech one. AI doesn't have values or judgment of its own. It reflects the choices of the people who build, deploy, and use it. That's why the real work happens in the human design choices behind it: what gets included and what gets left out, who gets consulted, what gets prioritized, and who's willing to ask the difficult questions before it's built.

From Challenges to Calls to Action

Every group closed its session with the same five-part template: name the problem worth solving, choose the SPHERES dimensions and what they mean for the challenge, identify what shifted, the key insight from the day, name the open question the region still needs to answer, and land on a concrete call to action. That structure produced every group's completed template featured under Key Themes.

Most groups defined "Human-Centered" and "Responsible" independently, for their own specific challenge, a structured early step in the WIN AI™ process. Their definitions are below.

Incentives for AI Adoption



Human-Centered: Meets human needs; AI adoption is a people problem, not a tech problem; rewards openness rather than penalizing non-use; prioritizes people over efficiencies and cost.

Responsible: Considers collective goals, impacts, and costs; AI aligned with values (with awareness leadership may hold outsized influence); prioritizes humans over financials.

Navigating Disruption



Human-Centered: Equal collaboration and accountability at all levels; attending to where and how disruption shows up; involving the disrupted in shaping the solution.

Responsible: Balancing AI's advantages against the disruption it causes; being accountable for that disruption.

Maintaining Human Control & Intention



Human-Centered: Building for humans' real needs from the start "baking it in," not "bolting it on."

Responsible Agility: Keeping a human genuinely in the loop while AI still drives value.

Pacing AI Adoption to Community Needs



Human-Centered: Preserving human agency so innovation happens when and where communities actually need it.

Responsible Agility: Keeping the community genuinely "in the loop" not just informed, but empowered to shape the pace.

WHAT SAN DIEGO IS ASKING FOR

CALLS TO ACTION

Synthesizing every conversation, panel discussion, breakout groups, and closing plenary into a set of concrete regional asks. These are not WE ARE IN AI™'s recommendations. They are what San Diego's community, working independently across five conversations converged on asking for.

INDIVIDUAL - Where agency starts

Start with yourself



Responsible AI use starts with individual choices: how you use it, what you accept, when you challenge it. Say the hard question out loud, someone else in the room may be sitting with the same one, and that's what starts the conversation.

POWER OF HUMAN AGENCY

Protect human creativity



Right now AI is bounded by human creativity. Human creativity is not bounded by AI. Protecting that starts close to home: show up for the artists and creative spaces in your own community.

POWER OF HUMAN AGENCY

Talk about what's being built



Understanding AI's consequences isn't optional anymore. Treat how to live a good life in the age of AI with the same seriousness as any technical question. Hold important conversations within your organization and your region about what's actually being built.

HUMAN-CENTERED INNOVATION

Exercise civic agency



Civic participation is AI governance: vote, attend budget meetings, engage with local AI policy, hold leaders accountable. California's new Engaged platform, launched to let residents weigh in on how AI affects their work and communities, is one place to start.

POWER OF HUMAN AGENCY

ORGANIZATIONAL - Where practice takes shape

Define healthy use



Involve employees and the community in AI's design, governance, and regulation, not after the fact, and make it someone's explicit responsibility to define and uphold what "healthy use" actually means.

RESPONSIBLE AGILITY

Include humans by design



AI should never be deployed as one-size-fits-all. Different people, different needs, different guardrails. That starts with baking human-centered values into the process from day one, not bolting them on after the fact, and keeping humans genuinely in the loop.

HUMAN-CENTERED INNOVATION

WHAT SAN DIEGO IS ASKING FOR

CALLS TO ACTION

ORGANIZATIONAL - Where practice takes shape

Require expressed uncertainty



AI systems should express uncertainty rather than produce confident-sounding wrong answers. Designed honesty over false precision.

EXPLICIT BY DESIGN

Invest in AI literacy



Help people understand how AI actually works: its probabilistic nature, how it ingests data, where it fails, and what it cannot know.

EXPLICIT BY DESIGN

REGIONAL - Where the region decides together

Pace adoption to need



Pace should follow community need, not technology availability. That takes sustained community conversation and ongoing engagement infrastructure. Citizens and regions need to keep asking "how much is too much?" , including the environmental cost of the infrastructure it runs on, before "how fast can we go?”

ENDURING IMPACT

Build practical guardrails



Comprehensive laws, stronger privacy protections, and clear accountability tied to data and policy, with someone whose explicit responsibility it is to define and uphold the standard.

ENDURING IMPACT

SECTORS REPRESENTED

VOICES FROM THE FORUM

Each group generated far more than any summary can carry. What follows is a second layer for additional details and questions pulled from the original group boards adding depth to the conversation.

BUSINESS Incentives for AI Adoption

The group mapped incentive design as a stakeholder problem before it was an economics problem. Business owners, employees, vendors, suppliers, and customers all named explicitly, with a separate note that a Chief Human Resources Officer (CHRO) specifically should carry outsized influence in these decisions. Working definitions of human-centered AI: it "allows people to learn and only use AI that benefits them," rewards openness to AI rather than penalizing non-use, and prioritizes people over efficiencies and cost. One note named the group's own starting point plainly: "organizations are at the beginning of AI initiatives." A separate thread asked the question underneath incentive design: **how do we understand money and power? That question paired directly with AI's utility question.** The group also observed that most small and mid-size businesses build their own AI tools rather than buy them, often without deep IT expertise, and a blind spot: "assuming AI is being deployed for the benefit of the company."

BUSINESS Navigating Disruption

Meeting those most affected where they are, at the forefront, that was the group's focus: a simple distinction between the people making decisions about AI, and the people living with the results of those decisions, and how easily the two groups can lose sight of each other. A hopeful counter-narrative surfaced too: not every disruption is a net loss. The group noted that disruption can also create new jobs that don't exist yet. The outcome, if this is done right: less stress, more creativity and innovation, **business and workforce both thriving, side by side**.

ECOSYSTEM Maintaining Human Control & Intention

For this group, it came down to one line: "**It's not about what AI provides**. It's about **how humans can best interact with AI**." Two questions stayed open at the end that didn't make it to the synthesis: how do we teach AI to de-escalate, and who should guide how we protect human intelligence?

COMMUNITY Preserving Humanness

This group wrestled with what it means to be human. One throughline was the idea of letting people evaluate their own essence, what makes them who they are, rather than letting a system do it for them. Two open questions carried real weight for the group: "AI doesn't have true emotions but emulates them, what emotions do we want it to learn?" and "what counts as healthy use of AI, and whose job is it to define it?"

Eight sectors were represented at the forum, including consulting, technology, higher education, workforce development, and community advocacy.

Sectors Represented, by Organization Count

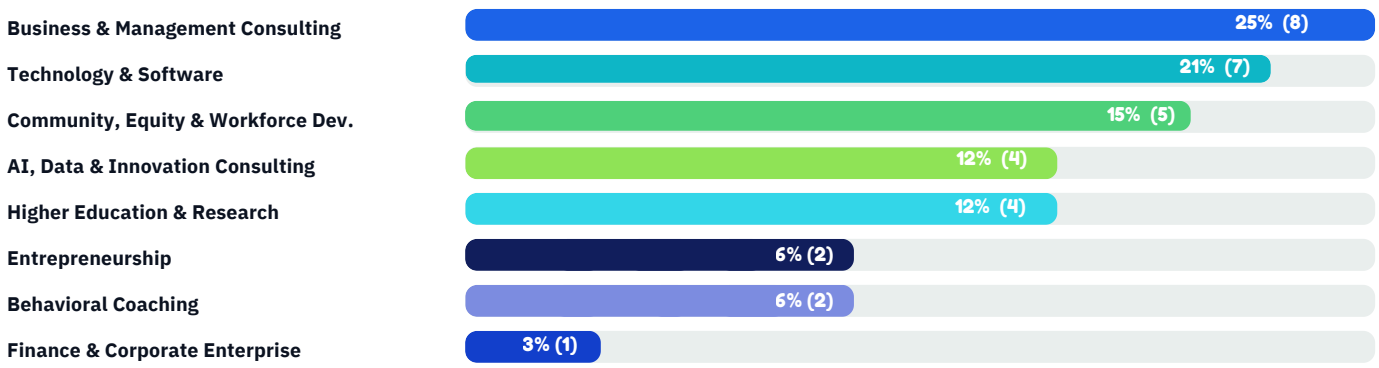

Counts reflect distinct organizations represented, not individual attendees. Some organizations represent more than one participant.

THE BIG PICTURE

THE CONVERGENCE



Five different starting questions. Working independently, San Diego's business, ecosystem, and community tracks, and the opening panel before them, converged on similar ideas. Participants questioned narratives we've stopped questioning, and the convergence across rooms that couldn't see each other's boards is the finding.

Live graphic recording captured during the closing plenary **"Looking Ahead."** Illustration by Liquid Sketch Studio (@LiquidSketchStudio) for WE ARE IN AI™.

Human-centered as the throughline

Every conversation defined it the same way: meeting real human needs, prioritizing people over pure efficiency or financial outcomes. This wasn't treated as a value statement. It was treated as a design constraint.

Shared accountability at every level

Responsibility cannot be delegated to leadership, to regulators, or to technology. It has to live at every level from individuals to executives deciding which processes to automate.

Agency as the antidote to fear

The antidote to disruption-driven fear isn't reassurance, it's participation. "Exercise whatever agency you have, whenever it shows up."

"This is a human problem, not an AI problem"

Multiple conversations reframed the challenge the same way: responsibility starts with how individuals choose to think about and use AI.

No one-size-fits-all

AI should not be designed or governed as if every person, community, or use case has the same needs. Personalization is not a feature, it's a requirement.

Everyone needs a say

Multiple groups independently rejected top-down rollout in favor of collaborative decision-making, representation across every function and level before anything gets decided.

Worth noting: the room skewed toward consulting and technology backgrounds, so this convergence, drawn from roughly 40 participants across 33 organizations, should be read as a strong signal from an engaged, self-selected group, not a statistically representative survey of the region.

REFERENCE MATERIAL

TOPICS CONSIDERED

Topics were proposed by individual post-it, then chosen by dot-voting. The highest-interest topics went on to become the day's five groups' themes. Naming often evolved from there, so the wording here won't always match a group's final theme, but it shows where the brainstorming actually started.

HIGH INTEREST

- Preserving humanity & co-creation
- Navigating this transition when disruption potential is so high
- Maintaining human control / agency / intention
- What are the limits? How to plan and think before AI
- Right incentives for people to adopt AI

MEDIUM INTEREST

- Reliability on AI for critical processes
- Helping small/mid businesses leverage AI the "right" way
- Taking care of our young people / education in the age of AI
- AI governance best practices in enterprise tech
- Business responsibility for AI's impact on society
- Balancing the tension between government & progress
- Reduce manual work & more time on what I enjoy
- Stress & Creativity
- Using renewable resources to power AI
- Defining tech that makes humans comfortable
- More comprehensive, representative training data
- Finding/scaling strength & being accurate with data

LOW INTEREST

- Equitable application of AI & accessibility
- What is the middle ground?
- AI in Education
- AI risk evaluation balancing innovation + safety
- AI's energy demands vs. climate change
- Teaching AI to say "I don't know"
- AI emulates emotions, doesn't have them

REFERENCE

METHODOLOGY

Every group ran the same WIN AI™ process, which carries an AI challenge from definition through prototyping, application, and continuous improvement, building the kind of shared, human-centered understanding that creates alignment, trust, and solutions that start with people, not technology.

The first three phases build the shared, human-centered vocabulary for AI that most organizations are missing. Before anyone starts building, literacy and upskilling work must come first. The full WIN AI™ process starts from there, carrying that same understanding all the way to a working, tested solution. Real innovation doesn't live in the technology; it lives in the people building it, in the intentionality they bring to what gets built, for whom, and why. A team fluent in that mindset builds something people actually need and are willing to trust.

W
Witness

Most teams go straight to building. This phase is different: groups build shared understanding of human-centered AI for their challenge, challenge assumptions, and focus on who it actually serves.

I
Intentionality

Each group selected two SPHERES dimensions, then used them to sharpen a broad topic into a specific, well-defined problem worth solving.

N
Navigate

Groups moved from understanding to action, articulating what shifted in their thinking, the open question their region still needs to answer, and a concrete call to action.

A
Architect

This phase is about building and testing an actual solution, translating the call to action into something ready for prototyping.*

I
Improve

An implementation isn't the finish line. This phase brings responsible agility to what comes next: feedback, human review, and the willingness to keep adjusting.*

*__Architect__ and __Improve__ are the next phases, the work of building and testing an actual solution. A single day is enough to define the right problem and reach a genuine call to action. The forum was designed to end at exactly that hand-off point. Each group's work ends with a defined problem, a chosen lens, and a concrete call to action, ready for what comes next.

The Process

One Structure, Five Stories

Before moving from a broad AI question to something more tangible, each group followed the same discipline built into the Witness phase: starting with a question rather than an answer, and sorting what the group knew, what it needed to know, and what it was only assuming. "Came in thinking, now believe, open question, call to action" carries that forward, building shared understanding before it ever reaches a call to action. That's by design. It's also what makes something else visible: self-organized groups who'd never worked together before, working through entirely different questions with no visibility into each other's boards, kept arriving at similar shifts and similar open questions. It's the shared structure that surfaced the convergence, five independent conversations arriving at the same place.

REFERENCE

METHODOLOGY

Definitions Built for the Challenge

Values and definitions are easy to agree on in the abstract, and that's exactly the problem. A value nobody had to define for themselves rarely survives contact with a real decision. That's the gap between believing in human-centered AI and actually practicing it, and it's exactly the gap this methodology addresses.

Human-centered design doesn't start with a definition handed down. It starts with the challenge itself: what it means for the people designing it, and what it means for the people who will live with what gets built.

Groups walked away with different definitions of the same two words. That's not a lack of consensus, it's proof the process worked. A definition a group builds for its own challenge does something a borrowed one can't: it gives everyone in the room a shared language for the problem, and a shared understanding of what they're actually trying to solve, before anyone starts building.

THE SPHERES

During the Intentionality phase of the WIN AI™ process, each of the breakout groups selected two of these seven dimensions as the analytical lens for its conversation, by group vote. SPHERES breaks human-centered AI into seven concrete dimensions, each one something a team can actually evaluate and act on.



System Resilience: Anticipatory design.
Build systems that adapt, recover, and stay centered on people as conditions evolve.

Power of Human Agency: Empowering by design. Chosen by 3 of 5 groups
Most AI systems are designed to make decisions easier, faster, or automatic. This dimension asks for AI that expands human capability instead.

Human-Centered Innovation: Build with people, not for them. Chosen by 2 of 5 groups
Empathy before architecture. The right problem before the solution. Systems designed to create value, not extract it.

Explicit by Design. AI is a probabilistic tool. Chosen by 1 of 5 groups
Design for that reality: help people understand uncertainty, question outputs, and make results more accurate over time.

Responsible Agility. Humans stay in the loop. Chosen by 2 of 5 groups
AI systems degrade. Models drift. Outputs shift. Responsible agility means feedback, accountability, and continuous improvement, by design.

Enduring Impact. Sustainable value. Chosen by 1 of 5 groups
Every efficiency gain is a distribution question: the human cost, who absorbs the cost, who saves time, whose role changes, and the environmental cost of powering the system that made the gain possible.

Strategic Alignment. A clear, principled north star. Chosen by 1 of 5 groups
Every AI initiative tied to a clear, principled north star. A strategy for AI is also a strategy for the kind of organization and society being built.

SPHERES is proprietary to WE ARE IN AI™. Selection counts reflect the two dimensions each group chose as its lens during the Intentionality phase of the WIN AI™ process.

AN INVITATION

What happened in the forum was inspiring. People walked in thinking about AI as a technology question and walked out understanding it as a human one. That shift is key. This forum was proof that it's possible, proof that a room full of community members, business leaders, technologists, educators, skeptics, and believers can sit down, ask fundamental questions, challenge assumptions, and leave with something they didn't have before: a shared starting point.

> **The trajectory of AI is not destiny. It is a series of decisions made every day, in the questions we choose to ask and the ones we choose to ignore.**
> **Designing the future starts with everyone.**

Spaces like this are how the future actually gets built, not by accident, but by people choosing to show up and think together before anyone starts building. The work doesn't end when the room empties. We need more spaces, and we hope what happened in this one is reason enough to take part in what comes next.

Taking part starts with a few real commitments:

1 Intention Over the Rush

Build with a long-term mindset, not the gold rush.

2 Build to Generate Value

Design systems that generate more value than they extract, for the people they touch.

3 Humanity at the Center

Technology doesn't decide what matters. People do. That's where real innovation actually lives.

That starting point is not the finish line. It is an invitation.

The future of AI should be shaped by everyone, not just the people building it or the people funding it. And the only way to get there is to start with people before we start with technology. Let the human-centered questions follow you into the spaces where AI is being discussed or implemented. Ask not just what AI can do, but why, and for whom. Exercise whatever agency you have. It is what turns a room full of people into a room that actually changes something, and it's what several groups at the forum, working independently, named as their call to action: exercise agency.

The same is true for organizations, at scale. Human-centered AI rarely happens from good intentions alone. It takes a repeatable process, clear ownership, and people upskilled to ask questions rather than lean on a tool to answer them. Systems tend to hold up better when the people they touch help shape them, not just hear about them once they're live. Not every problem needs AI in the first place, and knowing when it doesn't is part of the same discipline as knowing when it does. Every efficiency gain is still a distribution question, and the answer either becomes the foundation of something durable, or the explanation for why it wasn't.

Human-centered AI starts with the choices we make today. How this plays out doesn't come down to technology alone. It comes down to organizational leadership, and to the people the AI actually reaches.

ABOUT

“Design With People, Not For Them.”

WE ARE IN AI™ is a mission-driven organization and global ecosystem for Human-Centered Innovation and Responsible AI. Our work sits at the intersection of design, community, and technology. We develop frameworks, convene cross-sector conversations, and work toward AI adoption that is human-centered, because that is where real, enduring innovation lives.

San Diego's first Responsible & Human-Centered AI Forum is part of a broader program to build AI literacy and agency. Beyond the forum, we work with companies and public institutions to help people understand AI, find where it makes a useful difference, and put people at the center of how it is designed and adopted.

Human-centered AI shouldn't stay a conversation. It should be something teams and organizations can build, phase by phase, decision by decision.

The WE ARE IN AI™ Way

Human-Centered by Design
We design with people, not just for them. That shapes how we engage and how we work.

Create Value
Great organizations build sustainable innovation by creating value: people first, before profit.

Purpose Over Pace
At this inflection point, speed is not strategy. We start with questions and keep learning with the communities we serve.

Lifelong Curiosity
We explore AI and human-centered design with genuine openness.

Human Agency & Creativity
AI should expand human potential. Creativity and agency are not soft skills. They are the most powerful forces we have.

Enduring Impact
Not a feature added at the end. The foundation everything is built on.

We’re Just Getting Started

WE ARE IN AI™ will keep convening these spaces, in San Diego and beyond. The next Human-Centered AI Forum is already taking shape. Calls to action like these aren't the finish line for us. We're actively working to build the ecosystem and network that will help this conversation keep moving. If something here stayed with you, we want to hear from you at hello@weareinai.org, and we'd love to stay connected.

CLOSING

ACKNOWLEDGMENTS

Thank you to everyone who came together for this conversation. Our panel: Jonathan Behnke, Chief Information Officer of the City of San Diego; Mikey Knab, Business Outreach for the San Diego County Office of Labor Standards & Enforcement; and Di Le, AI Ethicist and Human-Centered AI Strategist. Special thanks to Di Le and Mikey Knab, beyond their contributions to the panel, for their editorial review of this document.

Thank you to every participant; we are grateful for your time and openness. The conversations and calls to action in this document are the product of our collective thinking, and its images are a reflection of community voices, captured in real time.

Thank you to our graphic recorder, Liquid Sketch Studio, who turned a day of conversation into something you can hold onto, and to AG video photo, for the event photography and video. We're also grateful to Kiln Rancho Bernardo for providing the space to host this conversation.

WE ARE IN AI™ convened this forum and produced this document. The organization was founded by Jarazet Altamirano.

The Nature of This Forum and the Pulse

The San Diego Responsible and Human-Centered AI Forum was an educational, collaborative experience, not an advisory session, and this document reflects that. The WIN AI™ methodology described here is presented for educational purposes only. The forum and its findings do not provide legal, financial, or regulatory advice. WE ARE IN AI™ is not responsible for independent decisions or actions taken by any party based on the findings of this forum or this document. The event was a public forum, and WE ARE IN AI™ is not responsible for what participants chose to disclose.

We've made every effort to ensure the information in this document is accurate, but full accuracy cannot be guaranteed. Readers should use their own judgment, and verify anything they intend to rely on. We reserve the right to modify the content of this document at any time.

This document may link to third-party websites and sources for informational purposes. WE ARE IN AI™ does not control these sites, cannot guarantee their accuracy, and is not responsible for their content. The findings, quotes, and characterizations of the forum's conversations reflect a synthesis of that day and are not necessarily the personal views of WE ARE IN AI™, its panelists, or its reviewers.

Synthesis & Editorial

The San Diego Human-Centered AI Pulse 2026 represents a synthesis of the conversations from participants. Participation in the forum does not imply endorsement of, or agreement with, the published takeaways. No individuals are named in the breakout group conversations; those are presented as a synthesis of discussion themes, not statements attributed to specific people.

Methodology

The WIN AI™ methodology and SPHERES used to facilitate and structure the forum and the session formats were developed and deployed by WE ARE IN AI™. These methodologies, frameworks, and tools may not be replicated, reproduced, adapted, distributed, or used in any form, in whole or in part, without the prior written permission of WE ARE IN AI™.

Participation in the forum does not grant any license or right to use, teach, or commercialize these methodologies. If you'd like to bring WIN AI™ to your own organization or region, we'd love to talk, reach us at hello@weareinai.org.

Terms of Use

This document is released under a Creative Commons Attribution–NonCommercial–NoDerivatives 4.0 license (CC BY-NC-ND 4.0).