

The Intelligence Is in the Interaction

Designing AI for How People Actually Work




Table of Contents

Introduction	3
Why AI Changes the Design Challenge	5
Working Intelligence, Failing Systems	7
The Same AI. A Different Outcome.	8
From Interaction to Design	10
Designing the AI Experience: Spotlight on AI Behavior	12
Eight Disciplines to Design Responsible AI	14
Five Questions to Answer Before You Build	21
From AI Capability to AI Experience	22
About Marlabs	23

Introduction

There's a lot of pressure right now to bring AI into every process.

Every leadership team wants AI. Every organization wants speed, efficiency, and measurable value.

But there's a larger challenge than whether you can build AI. The real challenge is understanding **where AI's value actually comes from.**

Most people assume it comes from the model. But the model is only part of the equation. **The intelligence you get from AI emerges through the interaction between the AI and the human using it.**

That's the central idea of this paper, and it's the most important piece that you as a business leader need to understand about AI. If you don't, you may find yourself wondering why **AI systems that perform well technically can produce disappointing results in the real world.**


One of the most important questions business leaders should ask themselves today is, "Can we design AI systems that people can trust, control, and use responsibly?"


AI that people don't trust will struggle to create value when they need to act on its outputs.

AI that people follow blindly creates risk.


The technology itself often isn't the root of the problems that ensue. It's the interaction you design around it.

We call this concept designing the space between the AI and the human.





**“The intelligence you
get from AI emerges
through the interaction
between the AI and
the human using it.”**



Why AI Changes the Design Challenge

Writing good software has always required good requirements. That hasn't changed.

What *has* changed is the role of the interaction.


Traditional software generally executes defined logic. But generative AI can create new outputs that were not explicitly defined in advance, including responses within a conversation.


A person asks a question, the AI produces an answer, the person adds context or challenges it, and the exchange continues. The value doesn't necessarily come from the first answer. It can emerge as the human and AI exchange information, add context, challenge assumptions, and refine the result.

Emergence of intelligence through interaction is not unique to AI. It happens throughout nature.

A single ant has limited capabilities. Yet, thousands of ants interacting with each other can solve problems, organize themselves, and even create structures such as bridges. These are all behaviors that no individual ant could accomplish alone. We see something similar when humans interact with one another. A conversation can produce an idea that neither person had when the conversation began.

AI creates another kind of interaction that we haven't experienced before.






You have a human being on one end and an interactive AI agent on the other that can search, retrieve, synthesize, and generate large amounts of information in natural language. The AI contributes information and the ability to generate responses. The human contributes context, judgment, experience, and accountability. **What results is more than either could contribute alone.**


This is why AI can be such a powerful augmentation tool.

Humans have a natural ability to augment themselves with technology. Some of the most primitive examples are tools, such as a sword or a staff. As we use the tool, our brains automatically incorporate it as an extension of ourselves, augmenting what we can do. AI is another form of augmentation. But it is an **intelligence augmentation tool**, and the value of that augmentation comes from the interaction.

That makes the interaction itself a business asset, and that's why leaders need to design AI intentionally.

"Traditional software generally executes defined logic. But generative AI can create new outputs that were not explicitly defined in advance, including responses within a conversation."





Working Intelligence, Failing Systems

Here's an example of that concept in action: imagine a global pharmaceutical company launches a new AI system that fills in medical records, drawing from patient information connected to the system.

The model performs well in testing, and the pilot is considered a success. But when the system goes live, trouble brews.

A specialist opens a patient record and sees that the AI system has suggested a medication name for the record. Something feels slightly wrong about the medication, but the AI recommended it, and the interface offers the specialist only two choices: **accept or cancel**. Accepting the medication is the default answer. It's faster. Canceling creates delays and extra work. The specialist accepts the recommendation.


But, in reality, the medication is wrong.


The failure wasn't caused by an incompetent clinician or a bad developer. It happened because the interaction made acceptance easier than questioning the AI. The system hid uncertainty instead of making it visible. The model generated an answer, but **it's the interaction with the human that ultimately shaped the outcome**.

Programmers will always miss requirements. Humans simply can't anticipate every scenario. That's why the difference between generative AI and traditional software matters and must be accounted for: generative AI can produce contextually relevant outputs even when the underlying information is incomplete, ambiguous, or uncertain.

When an AI encounters something it cannot resolve directly, it can still generate a plausible response. The result might look complete and confident, even though the underlying information is uncertain. In this case, the AI generated an incorrect medication name, but it provided no indication that the output required additional verification. The specialist saw an answer and assumed it was reliable.

The interaction turned an uncertain AI output into a human decision.






The Same AI. A Different Outcome.


Let's rewind that scenario. You have the same company using the same AI. This time, before launch, the team designed the interaction around uncertainty.

Specifically, the team defined how uncertainty should be detected and represented in the experience. When an AI answer is considered high risk, the system can trigger additional review, provide supporting information, require confirmation, route the decision to a qualified human, or document the decision without leaving the workflow.


Nothing about the underlying AI has changed. The interaction has changed.

The specialist is no longer deciding whether to trust a machine. She is working inside **a system designed to support human judgment**. That's the difference between an AI system that simply produces an answer and one that helps a human use that answer responsibly.





"Designing AI is about more than choosing a model or improving accuracy. It's about designing the space where people and AI work together."



From Interaction to Design

The pharmaceutical example shows what happens when an AI output enters a workflow without enough support for human judgment. But the same principle applies in much lower-risk situations.

Again, the interaction between AI and people is different from traditional software interaction, because AI synthesizes, analyzes, and generates massive amounts of information much more quickly than traditional software.


This enables humans to contribute more meaningful context, develop more meaningful insights, and craft more meaningful experiences **alongside the AI**.

Neither creates the final outcome alone.

Consider something as simple as identifying a weed growing rampant in your lawn. You take a picture and put it into two different AI models, asking what that weed is. Both AIs give you different answers. Both are incorrect. Yet, each AI responds confidently.

You might naturally interact further, explaining that you know the answer is incorrect and providing more context, such as your location, planting zone, or the fact that there's been unusually high rainfall.

As the conversation continues, a clearer, more accurate answer emerges. **The interaction changed the outcome.**





That's why designing AI is about more than choosing a model or improving accuracy. **It's about designing the space where people and AI work together.**

In an enterprise, the human-AI interaction extends far beyond what people see on the screen. It includes the workflows, information, business rules, controls, systems, and people that shape how the AI is used.

And that raises an important question for anyone designing an AI system.

If the intelligence emerges through the interaction, how do you design that interaction so that the right intelligence emerges?

That is the design challenge.

And the stakes are high. When AI acts or speaks for your business, the organization is ultimately accountable for how that capability is used and the consequences it creates.

The organizations best positioned to create value from AI will deliberately design the human-AI system around the interaction.

"When AI acts or speaks for your business, the organization is ultimately accountable for how that capability is used and the consequences it creates."

Designing the AI Experience: Spotlight on AI Behavior

The next step is turning these insights into a practical approach to create the conditions for people and AI to work together effectively.


What does it take to intentionally design how AI and people work together?

Building an effective AI experience requires you to make critical decisions before development about:

- Where AI enters a workflow
- How people interact with it
- What information people need to see
- How domain knowledge shapes AI behavior
- What happens when risk appears
- How the system is evaluated
- How the behavior is implemented
- Who remains accountable after launch


These decisions cannot be made by user experience (UX) teams alone. **AI behavior is a cross-functional delivery discipline.**

It defines how an AI system responds to people, information, uncertainty, risk, business rules, and changing conditions within a business workflow.



Defining AI behavior brings UX, engineering, business, data, process, risk, compliance, and operations together to establish how the AI system should behave before those decisions become embedded in the product.

The approach can be summarized in five steps:





This is where UI/UX design creates value by making that behavior visible, usable, and actionable in the workflow.

Eight Disciplines to Design Responsible AI

No single discipline or function can define AI behavior alone. Responsible AI behavior is created when the people, processes, workflows, data, business, compliance requirements, engineering, operations, and UX perspectives come together before development begins.

This alignment is the foundation for the rest of the process. As part of the Marlabs AIx framework for UI/UX design, the following eight disciplines work together to define how an AI system should behave and how people should interact with it.





1. Human Interaction Design

Human interaction itself is one of the most important aspects of AI behavior to address. Your design should support trust without encouraging blind acceptance.

People need to understand what the AI is doing and what its output means. They need to know when they can act on an answer, when they should question it, and what they can do when they disagree with the output, such as reject or escalate it.



Addressing these factors before build helps create a human-AI interaction that is trustworthy, usable, transparent, actionable, and controllable.


2. Process and Workflow Design

AI operates within a workflow, not in isolation. Therefore, it's important to define where AI enters the process, what decisions or actions it is allowed to take, where it should pause, and where human checkpoints are needed for critical or sensitive decisions.

A high-risk recommendation, for example, might require a human-approval checkpoint. An uncertain result might need to pause the workflow rather than move directly to completion.

In short, process and workflow design answers, **"Where should AI act, where should the human act, and how should they work together?"**





3. Information Flow

People cannot make informed decisions about AI output if the information they need is hidden. The interaction may need to make relevant data, uncertainty, rationale, sources, or status visible to the user.

This doesn't mean showing people everything the system knows. It means designing what the person needs to see to understand the AI's output.


The information available to a person becomes especially important when the AI is uncertain. When considering information flow, a good question to answer is, **“What does the person need to see to be able to decide what to do next with confidence?”**


4. Knowledge Design

AI behavior is shaped by the knowledge surrounding it. Domain knowledge, rules, thresholds, exceptions, and context all influence how an AI system should behave within a particular business environment. Knowledge design defines how all those factors shape AI behavior.

An AI system operating in a clinical environment will have different requirements from one supporting a customer service workflow or an internal business process. That's why the AI experience should be designed around the business context in which the AI is being used.

That context should also determine whether AI is appropriate in the first place, instead of assuming that AI belongs in every step of the process.






5. Risk and Compliance



AI governance has to become part of the interaction to turn abstract risk into concrete human action. Define what needs to be reviewed, documented, audited, restricted, or controlled. This includes privacy, security, auditability, and human oversight.

Here are some examples.

- If the system detects uncertainty in a high-risk output, the defined behavior might be to trigger a review-required state.
- If data is missing, it might result in a missing-data signal.
- If a field is high risk, the defined behavior might require a human-approval checkpoint.
- If an output is unclear, the user might need a challenge or escalation path.
- If a regulatory requirement calls for documentation, the interaction will need to support traceability and create an appropriate audit trail.
- If a request is outside of the system's intended scope, the AI may need to refuse or provide a fallback.

These are behavioral decisions. The interaction makes them visible and actionable in the workflow.





6. Evaluation and Monitoring

The AI experience doesn't end when the system launches. Evaluation and monitoring determine how well the system behaves as intended after it goes live and people use it in real workflows.


In this part of the design, you determine how behavior is tested, measured, and improved after launch.


Evaluation should answer questions such as:

- Are people using the AI appropriately?
- Are they challenging outputs when they should?
- Are certain behaviors producing unexpected results?
- Is the system continuing to perform as intended?

Monitoring then helps identify changes or emerging issues over time, such as shifts in usage, unexpected behaviors, declining performance, or new risks. The answers can reveal where the AI behavior or interaction needs to change.

These measures should ultimately connect to business outcomes, such as reduced cycle times, increased productivity, improved quality, better customer experience, revenue growth, or reduced risk.





7. Engineering

Once the behavior has been defined, the engineering team translates those decisions into the system. The team defines how the agreed behavior becomes system logic, states, APIs, permissions, constraints, and implementation.

Engineers shouldn't have to figure out these behavioral decisions on their own while building the system. They need to be part of the behavioral design from the beginning so that the intended behaviors are defined before they become code.



8. Business Ownership


AI participates in decision making, but when AI acts or speaks on behalf of your business, the organization remains accountable for how that capability is used and the consequences it creates.

Business ownership is part of the AI experience itself. It shouldn't be addressed after the product is built. Business leaders need to establish ownership during the design process.


Define:


- Who owns the decision?
- Who owns the risk?
- Who owns adoption?
- Who is accountable for how the system behaves after launch?





"Responsible AI behavior is created when the people, processes, workflows, data, business, compliance requirements, engineering, operations, and UX perspectives come together before development begins."





Five Questions to Answer Before You Build

Once you align those disciplines, the next step is to answer five questions before writing a single line of code.

**Need:**

Do we actually need AI, or is there a simpler solution?

**Data:**

Is the data ready, reliable, and appropriate for the task before we start? Data readiness includes the quality, completeness, accessibility, security, and governance of the data the AI depends on.

**Uncertainty:**

Who decides what the AI does when it is uncertain? Who reviews, challenges, or escalates?

**Evaluation:**

How do we evaluate and monitor the AI system after launch to know it is working as intended?


**Ownership:**


Who is accountable for the AI's behavior after launch?

Most teams answer these questions too late and run into trouble.

If no one decides what happens when the AI is uncertain, the system will have no defined behavior when uncertainty occurs. If no one defines how a user should challenge an answer, the user will still have to decide what to do. If no one establishes who is accountable, accountability will still be needed if something goes wrong.

Answering these questions thoughtfully from the start protects you from having to fix preventable production issues, user confusion, support tickets, compliance concerns, or failed adoption later.





From AI Capability to AI Experience

The five questions really show that AI should not be treated as a capability that is simply added to a process. It must be designed as an experience in which people and AI work together.

An AI capability gives an organization access to what the technology can do. But an AI experience determines how people use that capability, how the workflow responds to it, and how the organization manages it. That distinction matters because an AI model can be technically capable and still produce a poor business outcome.


The value goes beyond making AI more accurate or more powerful. You need to create the conditions for people and AI to produce better outcomes together. That requires organizations to define AI behavior before it becomes embedded in the product:


- Where AI acts
- Where people act
- What information people need
- What happens when the system is uncertain
- When human judgment is required
- Who remains accountable

This is the purpose of Marlabs' Alx framework for UI/UX design within the broader **AgilityAI** framework. Alx brings the necessary disciplines together to define AI behavior, shape the human-AI interaction, design the workflow and controls, and establish how the experience will be evaluated after launch.

The objective is to design the space between the AI and the human so that the intelligence that emerges from that interaction is useful, trusted, controllable, and accountable.

Ultimately, the value of AI is not determined by what the model can do on its own but what people and AI can accomplish together.





About Marlabs

Marlabs is a global AI and digital transformation partner that helps enterprises move from AI experimentation to measurable business impact. We specialize in agentic AI implementation, workflow redesign, enterprise transformation, intelligent automation, and building responsible human-AI operating models.

We can help you redesign workflows, reduce friction, and build effective governance so that your organization can unlock sustained value from a human-AI workforce.

Ready to Take the Next Step?

Ready to move beyond AI pilots and redesign how your teams execute their work? Marlabs can help you reimagine workflows for AI transformation, reduce operational friction, and build scalable agentic operating models.

[Contact us](#) to schedule a discussion.

Learn more about our AI services:

[AgilityAI](#)

[AI Consulting Services](#)

[AI Transformation](#)

"In the agentic era, organizations that can repeatedly re-architect how people and AI work together will gain a lasting competitive advantage. This work must begin now."

 marlabs
driving digital agility

