

A FOUNDING MANIFESTO

**Better ways to think.
Better ways to build.
Better ways to live.**

A FOUNDING MANIFESTO

Our purpose is to find better ways to solve problems, make intelligence more accessible, and help people become more capable.

This document sets out the principles behind yukti, the kind of research we want to pursue, and the commitments we intend to work towards.

It is a starting point for us, not a claim that we already have all the answers.

Contents


01	Why yukti Exists
02	The Right Way to Build Intelligence
03	The Internet Changed Access. AI Can Take It Further.
04	Workhorses for Everyone, Frontier Models for What Demands Them
05	AI Should Make You Better, Not Mindless
06	From AI Users to AI Researchers
07	The Future of Work, Value and Human Agency
08	Safety, Responsibility and Access
09	Intelligence Has a Cost
10	What We Are Building
11	Our Commitments
12	A Final Word

01

Why yukti Exists

There is almost always a better way to solve a problem.

A simpler way. A more resource-efficient way. A way that produces better results without requiring unnecessary money, time, energy or complexity.




Finding that way is yukti.

In Sanskrit, yukti carries the sense of reasoning, ingenuity, and the right method of fitting the means to the end. It is not simply about having intelligence. It is about knowing how to use it.

That distinction matters.

The world does not need intelligence for the sake of intelligence alone. It needs intelligence that solves real problems, improves lives, and makes things possible that were previously out of reach.



We are building our research lab around this belief.

We want to build powerful models, efficient systems, useful tools and better ways for AI to learn. We want to bring together capability, cost, safety and human priorities instead of treating them as separate concerns.

We believe the best solution is not always the biggest model, the longest answer, the most complicated architecture or the most expensive system.

Sometimes, a smaller model with the right tools and memory is better. Sometimes, a specialised system is better than a general one. Sometimes, the right answer does not require generating paragraphs of text at all.

The challenge is to discover what works best for the problem at hand.


We are not here to make intelligence bigger at any cost. We are here to make intelligence work better.

02

The Right Way to Build Intelligence

Much of the AI industry has been driven by a race for greater capability. This has produced extraordinary advances, and we want to push those advances further.

But capability is not the only measure of progress.



A model can be impressive in a research paper and still be too expensive, unreliable or difficult to use for an ordinary business. A system can produce more output without producing more value. A tool can save someone an hour today while making them less capable of doing the work tomorrow.

We should ask more than, "How powerful can we make this model?"

We should also ask:

- What are people actually trying to accomplish?
- What level of quality do they need?
- What are they willing and able to pay?
- What risks come with the system?
- Can we achieve the same result with less compute, money and energy?
- Does using this system leave people better equipped to solve the next problem?



These questions should shape research from the beginning, not be added after a model has already been built.

This does not mean choosing the cheapest option every time. Sometimes, spending more is entirely justified. Cancer research, for example, may warrant substantial resources because the potential benefit is enormous.

The right decision depends on the problem, the expected benefit, the cost and the risks.

We want to make those trade-offs deliberately.


The goal is not minimum capability or minimum cost. It is the right capability, at the right cost, for the right purpose.

03

The Internet Changed Access. AI Can Take It Further.

We often talk about AI as if it were the first time humanity gained easy access to intelligence. It isn't.

Humanity has spent centuries finding ways to make knowledge, expertise and intelligence more accessible. Libraries, schools, universities, books and the internet have all expanded our ability to learn and solve problems.




The internet did not make libraries obsolete. Libraries still matter. What the internet changed was the cost, speed and reach of finding information. Knowledge that once required access to a particular building, institution or person could increasingly be reached from almost anywhere.

It also accelerated what people could do with that knowledge. A student could learn from material produced on the other side of the world. A developer could find documentation, study other people's code and build something new without having to know the right person first.

The underlying expertise did not disappear. Access to it changed.

AI can take this further.



Information is not the same as understanding, and access to information is not the same as having someone help you use it. AI can help explain unfamiliar concepts, connect ideas, explore alternatives and guide someone through a difficult task. It can help a person move from having a question to investigating it, and from understanding an idea to putting it into practice.

Imagine someone with no formal background in AI research. They have a question about how a model learns, but they do not know where to begin. With the right AI tools, they can learn the foundations, examine existing research, test an idea and gradually develop the skills to contribute something of their own.

The person still has to think. The research still has to be checked. Expertise still matters. But the starting point has changed.

That is the kind of expansion of human capability we care about.

We want to make the benefits of intelligence available to more people, regardless of where they live, what resources they have or whether they began with the right connections.


The opportunity is not simply to give everyone answers. It is to give more people the ability to understand, discover and do things that were previously beyond their reach.

04

Workhorses for Everyone, Frontier Models for What Demands Them

We believe in powerful intelligence.

Frontier models are extraordinary achievements. They can help us understand difficult problems, conduct research and discover approaches that would otherwise remain beyond our reach. We want to learn from them, work with them and eventually contribute to advancing them.




But most people do not need frontier-level intelligence for every task.

Most people need good workhorses.

They need dependable AI that helps them write code, understand unfamiliar subjects, analyse information, run a business, learn a skill or complete everyday work. They need something they can depend on without worrying that every useful task will become prohibitively expensive.

We want to build those workhorses.

As research advances, capabilities that once required the most advanced models gradually become available in smaller, more affordable systems. What is exceptional today can become ordinary tomorrow. We want to help accelerate that transition.



At the same time, some problems genuinely require more powerful systems. Those systems should exist, and people should be able to use them when the task justifies the additional cost and responsibility.

Think about personal defence. Everyone should have the opportunity to learn how to defend themselves. Everyone should have access to the knowledge and training that make them more capable of protecting themselves.

But that does not mean everyone needs a tank.

Not everyone needs a tank. But everyone deserves a good workhorse.

The point is not that powerful tools should be reserved for a privileged few. It is that different capabilities serve different purposes, and the systems using them must be appropriate to the risks involved.


We want a world in which powerful intelligence exists, practical intelligence is affordable, and people can choose the level of capability their work actually requires.

05

AI Should Make You Better, Not Mindless

Productivity is not enough if it comes at the expense of human capability.

Imagine using AI to become a better developer. You learn unfamiliar concepts faster, explore more solutions, understand why your code fails and build things you could not have built before.



Now imagine a different outcome. You produce more code, but gradually lose the ability to understand, debug or write it yourself.

Both outcomes might look productive on a dashboard. They are not the same.

We want the first.

AI should help people learn, reason, explore and create. It should let a person with no formal research background investigate a scientific question. It should help an aspiring developer become a stronger developer. It should let a curious person pursue an idea that previously required an entire team.

AI can be an extraordinary partner in thinking. It can juggle ideas, connect concepts and help people explore possibilities at a speed that would otherwise be difficult to achieve.



People should be able to become mad researchers with AI, following their curiosity further than ever before.

But they should not become mindless in the process.

Physical technology reduced the amount of physical effort required for many tasks. We eventually built gyms because keeping our bodies strong still mattered.

As AI takes on more intellectual work, we should take seriously the question of how we keep our minds sharp, too. We may eventually need the equivalent of brain gyms: ways to practise, challenge and strengthen our thinking while benefiting from intelligent tools.

We do not yet know every answer. But the question must be part of how we build.

Our systems should help people understand, not merely receive. They should help people develop judgment, not merely outsource it. They should make people more capable of doing meaningful work, not just more dependent on a machine to do it.


AI should expand human capability without needlessly diminishing the capabilities that make us human.

06

From AI Users to AI Researchers

AI systems improve through research, data, evaluation, feedback and iteration. But there is another source of improvement that deserves much more attention: the people who use these systems every day.

Consider a familiar interaction.



You give an AI a task. It fails. You explain what it misunderstood, correct its approach and steer it towards a better result. The next attempt succeeds.

That loop happens again and again.

A user discovers an edge case. A developer finds a better way to solve a programming problem. A financial professional identifies a mistake in an analysis. Someone working in law, healthcare or supply chains finds that the system lacks a capability their work requires.


Across millions of interactions, people develop knowledge about what AI gets wrong, how it can improve and what better performance looks like.

Some of those discoveries are valuable research contributions.

Today, much of that value can disappear into a conversation. We believe there should be a better way.

We want to create a system in which people can submit meaningful contributions, have them evaluated, and receive recognition when their work demonstrably improves our systems.

The principle is straightforward:



If your contribution helps create value, you should have a path to share in that value.

This is not a promise that every correction or conversation will earn money. Contributions must be evaluated. Some will be useful; others will not. We must be able to identify what improved, attribute contributions fairly and understand which systems benefit.

But when someone makes a meaningful contribution, we want the relationship to extend beyond a thank-you message.

Our long-term ambition is to connect models, capabilities, evaluations and products so that the value created by a contribution can be traced more meaningfully. A contribution that improves a capability used repeatedly should not necessarily be treated as if it had value only once.

We also want to create a path for exceptional contributors to become researchers.

That path could begin with recognition and named credit, continue through paid research work or a focused research collaboration, and eventually lead to a full-time role for those who want one. People with deep expertise in a particular field should be able to help shape the intelligence used in that field.

You should not need to begin your career at an AI lab to have a chance to contribute to AI research.

We want to build open-source tools and frameworks that allow other AI labs to create similar systems in their own way. We believe broader participation can make AI development more useful, more diverse and more connected to the real world.


This is how we want to approach the future of research: not as a closed activity conducted only by a small group of experts, but as a field in which more people can participate meaningfully.

07

The Future of Work, Value and Human Agency

We do not know exactly what the world will look like when AI can perform a much larger share of economic work.

New industries may emerge. Entirely new forms of work may become possible. Entertainment, games, films, tourism and experiences we cannot yet imagine may become more important. Some existing roles will change, and some may disappear.




We should be honest about that uncertainty.

But one thing should remain clear: the economy exists for people. People do not exist for the economy.

If AI eventually performs much of the work that currently gives people an income, society will need to think seriously about how people share in the prosperity that AI creates. Universal basic income may be part of that future. We should not assume it is the only answer, nor that it alone will preserve people's autonomy.

Security matters. So do goals, growth, contribution and the freedom to pursue something meaningful.

People should have the option to live without constantly having to prove their economic worth. They should also have opportunities to learn, build, research, create and participate in the value they help generate.



We see a possible future in which people use AI to complete work, exercise judgment and take responsibility for outcomes.

Consider an AI-generated poster. The machine can produce the image. But someone still has to decide whether it is good, whether it serves the purpose, whether it should be published and whether they are willing to stand behind it.

As execution becomes easier, taste, ambition, accountability and judgment may become increasingly important.

Some people will use AI to do their own work. Others will use it to deliver completed work to people who want the outcome without managing the process themselves. Humans will continue to shape what is worth doing, what good looks like and what should be done at all.

We want AI to increase people's agency, not make their lives revolve around the decisions of a handful of companies.


The future should give people more ways to participate, more opportunities to grow and more control over what they choose to pursue.

08

Safety, Responsibility and Access

We believe powerful intelligence should be developed responsibly, and that responsibility must grow with capability.

Our position is not that powerful AI should never exist. Nor is it that every person should automatically have access to every capability, regardless of the risks.



Think again about the stick and the tank.


A stick can be used for self-defence, to support someone while walking, or for many ordinary purposes. It can still hurt someone, so context and responsibility matter. But the risks it creates are generally different from those created by a tank.

A tank can cause vastly greater damage. Making it available responsibly requires far more than handing someone the keys. It requires accountable institutions, trained operators, clear rules, appropriate oversight and systems that can answer for how it is used.

AI capabilities differ in a similar way.

A workhorse model that helps someone draft an email is not the same as a highly capable system that can independently execute complex actions with significant real-world consequences.

We should make useful intelligence widely accessible. But when capabilities can cause greater harm, their deployment should require safeguards and accountability appropriate to what they can do.



Workhorses for everyone. Frontier models for what demands them, within systems capable of handling them responsibly.

This principle should shape how we think about model capability, tool access, autonomy and deployment. The safeguards should reflect the actual risks of a system, rather than treating every AI application as equally dangerous or equally harmless.

Responsibility is shared, but it is not interchangeable. AI developers are responsible for the systems they build and the safeguards they provide. Organisations deploying AI are responsible for how they integrate it into their operations. Users are responsible for their own decisions and actions, within the limits of their role and control. Responsibility should follow what each party can reasonably understand, influence and prevent.

We also believe that progress should not stop simply because a technology is powerful. The answer is to understand risks, build appropriate controls, evaluate systems carefully and make decisions that reflect the context in which they will be used.

Our ethical foundation is informed by dharma (धर्म): the idea that power and action carry duties, that means matter as well as ends, and that decisions must be considered in light of their consequences for others.

We want to ask not only whether something can be built, but whether it should be built, under what conditions, and for whose benefit.

Our concern extends beyond immediate human convenience. We should consider the welfare of other living beings, the communities affected by our decisions and the world that future generations will inherit.

Safety is not a feature to bolt on once a model becomes powerful. It must inform the way we research, build, evaluate and deploy intelligence from the beginning.

09

Intelligence Has a Cost

Intelligence does not exist in a vacuum.

Models require computation. Computation requires hardware, electricity, infrastructure and other physical resources. Data centres have real costs and real consequences.





FIG. 17

TWO MACHINES, ONE CUBE

We should not treat those resources as if they were infinite.

Our ambition is to get more useful intelligence out of every unit of compute, energy and money we use. This means improving models, developing better methods, using specialised systems when appropriate and designing software that avoids unnecessary work.

Efficiency is not merely a way to lower our bills. It is a way to make useful intelligence accessible to more people while reducing unnecessary resource consumption.



We also want AI to help solve problems that matter beyond the technology industry: water scarcity, pollution, scientific challenges and other problems that demand better use of our collective intelligence.

There will be cases where the benefits of a system justify substantial resource use. There will also be cases where a simpler solution is better.

We should make those decisions with care.


The goal is not to consume the most resources in pursuit of intelligence. It is to use intelligence to make better use of our resources.

10

What We Are Building

yukti is an AI research lab. Our products are not separate from our research; they are places where our ideas meet real problems.

Our current research direction has several connected ambitions.



01 Institutional memory

Organisations accumulate enormous amounts of knowledge across documents, conversations, decisions and people. Much of it is difficult to retrieve or use when it matters.

We want to build systems that help organisations remember what they know, investigate questions deeply and make better use of their collective knowledge.

02 Personal memory

People should be able to work with AI over time without repeatedly starting from zero. We want to explore persistent, useful memory that helps AI understand a person's ongoing work and context.

03 Workhorse models

We want to build capable, dependable and cost-effective models that people can use for everyday work. The goal is to bring useful intelligence within reach of more people without sacrificing the quality they actually need.

04 Better ways for models to learn

We want to research methods that help models improve, generalise and perform better on specific tasks. Different users have different needs, and improvements should not have to follow a one-size-fits-all path.

We want to make it easier for people to define what good performance means for their work, create their own benchmarks and help systems improve against them.

05 An automatic research lab for every company

Choosing an AI model should not require guessing which model, tools or architecture might work best.

We envision systems that automatically compare models, harnesses and tools for a particular use case, evaluating quality, cost and safety to identify the best approach.

Instead of asking every business to become an AI research team just to choose the right system, we want to make that research available to them.

06 Safety and alignment

Safety is not a separate destination that begins after capability research is complete. It is a continuing responsibility that shapes our other work.

We want to investigate how capable systems can remain useful, reliable and appropriately controlled as their abilities grow.

07 Embodied and physical intelligence

Intelligence will not remain confined to text and software. We are interested in what becomes possible when AI interacts with the physical world through robotics and other systems.

This is a longer-term research ambition, not a promise about a particular launch date. We want to approach it with the same attention to usefulness, resource efficiency, safety and human benefit.

These areas will influence one another. Research into learning can improve workhorse models. Better memory can improve research agents. Safety findings can change how every product is built.


This is our current understanding of where the most important opportunities lie. It is not a rigid sequence of milestones, and it is not a claim that we already know the best way to pursue every one of them.

11

Our Commitments

A founding philosophy matters only if it changes how an organisation behaves.

The following commitments guide the work we are doing now and the systems we intend to build over time.



01 Build for real use

We will judge our systems by how well they solve real problems, not only by how impressive they appear in demonstrations. We will use benchmarks, evaluations and practical testing to understand where they succeed and fail.

02 Treat cost as part of capability

We will work to reduce unnecessary computation and expense. Where a smaller or more specialised system can do the job well, we will take that possibility seriously.

03 Help people grow

We will aim to build systems that help people understand, learn and develop their abilities, rather than measuring success only through the amount of work produced.

04 Make safety a pre-requisite not an afterthought

We will consider risks, limitations and safeguards throughout research and development. We will not treat capability as a substitute for responsibility.

05 Create a path for contributors




FIG. 21

LEARNING IN PUBLIC

As our products and resources mature, we intend to build a programme that recognises and rewards meaningful contributions to our systems.

Our initial ambition is to establish a defined pool targeting 2–5% of an agreed revenue or profit base for validated contributions, subject to the programme's published terms, financial viability and implementation.

The precise calculation, eligible products, attribution rules and payment conditions will be established before the programme launches. We will need a fair way to measure which contributions improve performance, how those contributions benefit different products, and how the resulting value should be shared.

This is a future commitment, not a claim that such a programme is already operational.

We also intend to develop pathways from contribution to paid research, collaborative projects and potential full-time opportunities.

06 **Share the tools**

Where practical, we intend to open-source tools and frameworks that help other organisations evaluate AI, improve learning systems and build contribution-based research programmes.

We want the benefits of these ideas to extend beyond yukti.

07 **Keep learning in public**

We want our research claims to be supported by evidence. As our systems mature, we intend to publish benchmarks, explain our methods and communicate limitations honestly. We will not present an unvalidated ambition as an achieved result.

These commitments will develop as we gain experience, but we should explain material changes rather than quietly rewriting our promises.

End

A FINAL WORD

A Final Word

We do not believe we can predict every consequence of artificial intelligence.

We do believe we can choose the principles that guide our work.




FIG. 22

INTELLIGENCE, GROUNDED IN THE WORLD

We want to build powerful intelligence without confusing power with progress. We want to make useful systems affordable without assuming that cost is the only thing that matters. We want to advance research without losing sight of safety, sustainability or the people who will live with its consequences.

We want a world where more people can learn, create, investigate and solve difficult problems. Where people can contribute to the systems they use. Where meaningful contributions can lead to recognition, opportunity and a share in the value they help create.

We want AI to help people become better at being human, not less involved in their own lives.

The name yukti reminds us that the right answer is not always the most obvious one. Often, the real work lies in finding a better method, joining the right pieces together and understanding what the problem actually demands.

That is the kind of intelligence we want to build.

Not intelligence detached from the world, but intelligence grounded in it.

Not capability without responsibility, or productivity without growth.

Intelligence that helps people do more, understand more and become more capable of shaping their own future.

| That is yukti.

Vaaron Sinha

VAARUN SINHA
FOUNDER, YUKTI



USEYUKTI.COM

