



TAKSHASHILA
INSTITUTION

AI Adoption - Think Tasks, not Jobs

A framework for AI adoption in India's services sector

Arindam Amaresh Goswami, Sridhar Krishna, Tara Chawla

Takshashila Discussion Document 2024-22
Version 1.0, November, 2024

This document outlines the capabilities of AI with respect to jobs and proposes frameworks to consider which tasks will be replaced by AI, with a particular focus on services.

Recommended Citation:

Arindam Amaresh Goswami, Sridhar Krishna, Tara Chawla, "AI Adoption – Think Tasks, not Jobs", Takshashila Discussion Document 2024-22, November, 2024, The Takshashila Institution.

Executive Summary

This discussion document examines the potential impact of artificial intelligence (AI) on employment, with a particular focus on services. It argues that while AI is highly likely to transform employment in all sectors, especially in services, when adopted responsibly, it's unlikely to cause mass unemployment in the near future. AI's adoption in real-world settings is often slower due to factors such as implementation costs, process changes, and risk assessment. The paper proposes a framework for understanding how tasks may transition to AI, considering factors like technology maturity, return on investment (ROI), potential harm from errors, and AI explainability. The research emphasises that while some roles may be replaced, new positions may emerge to manage and oversee AI systems.

The document stresses the importance of upskilling and reskilling to adapt to an AI-driven economy. It highlights potential barriers to AI adoption, including resistance from labour unions and industry lobbies. The authors recommend that policymakers focus on helping people acquire skills necessary for an optimal transition to an AI-driven economy, and suggest that industry should focus on transitioning a significant portion of its workforce to more AI-complementary roles to maintain its competitive edge.

This document has been formatted to be read conveniently on screens with landscape aspect ratios. Please print only if absolutely necessary.

Authors

Arindam Amaresh Goswami is a Research Analyst working in the High-Tech Geopolitics Programme at The Takshashila Institution.

Sridhar Krishna is a Senior Scholar at The Takshashila Institution. Prior to joining Takshashila, he was a Managing Director at Accenture's AI practice.

Tara Chawla worked as an intern at The Takshashila Institution.

Acknowledgments

The authors would like to thank Pranay Kotasthane for his valuable comments and feedback.

Table of Contents

Executive Summary.....	2
Table of Contents	3
Introduction.....	5
1. What Can AI Do Better Than Humans?	9
2. What technologies are Driving AI Adoption?	10
3. Distribution of Employment Across Sectors	14
4. AI Capability vs AI Adoption.....	16
5. Proposed Framework for Transition to AI	18
5.1 What AI can do?.....	19
5.2 Return on investment vs AI technology maturity	20
5.3 Harm vs explainability	24
5.4 Other implementation barriers.....	26
5.5 AI deployment.....	27
6. Deepdive into Information Technology Sector.....	28

7. Deepdive into Education Sector37

Conclusion43

References45


Introduction

There is growing evidence that artificial intelligence is likely to be able to do many things humans can do and do some tasks even better than humans, as illustrated by this chart below from Our World in Data. “The language and image recognition capabilities of AI systems have developed rapidly.”¹

Test scores of AI systems on various capabilities relative to human performance

Our World
in Data

Within each domain, the initial performance of the AI is set to -100. Human performance is used as a baseline, set to zero. When the AI's performance crosses the zero line, it scored more points than humans.



Data source: Kiela et al. (2023)

OurWorldInData.org/artificial-intelligence | CC BY

Note: For each capability, the first year always shows a baseline of -100, even if better performance was recorded later that year.

Figure 1: AI Capabilities vs Human Performance¹

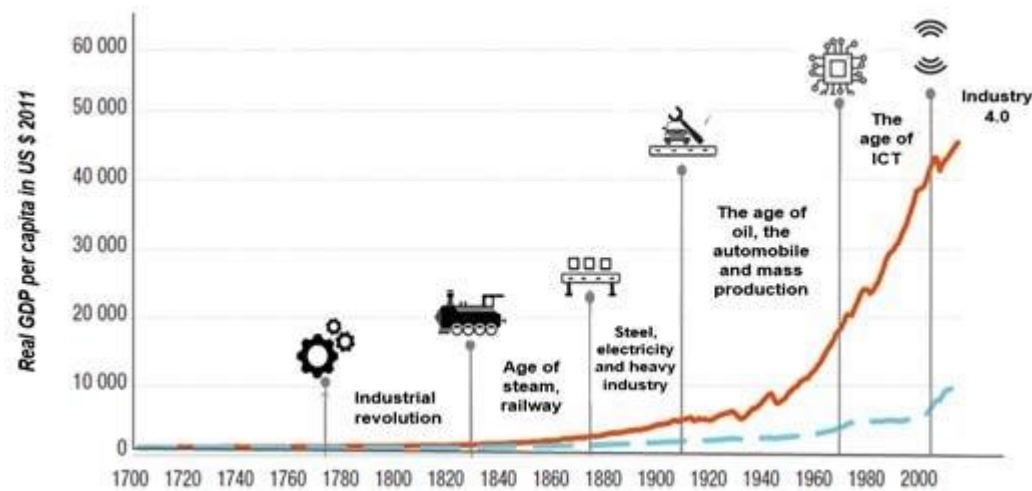
Will Artificial Intelligence replace humans in everything humans do today and create mass unemployment? This document seeks to understand what AI can do today, the upcoming advancements in technology and how we expect jobs to change.

At the outset, we want to point out that AI, like any other technology, does not take over jobs directly. It takes over tasks. Let's think of jobs as responsibilities people are given for a particular outcome and tasks as the way people carry out their jobs. In some cases, it may mean some jobs disappear because AI can deliver the outcome on its own, but the path to knowing this is to analyse the tasks involved and how using AI is likely to be beneficial.

We shall start by looking at the state of AI technology, understand what AI can do better than humans and then evaluate the constraints to AI adoption.

It is useful to step back a bit and look at some lessons from history before we examine AI and its effects on the economy and jobs.

The chart below shows how the global economy exploded after the industrial revolutions in the past.

Figure 2: Growth of Global GDP²

Global GDP grew at between 0.2% and 0.3% for over a 1000 years but grew relatively exponentially after the industrial revolutions.

Each of the earlier industrial revolutions actually generated more jobs than there were before.


When other technologies like mechanisation, electricity, and computers came into being, society worried, like they do today, about whether the new jobs created will exceed the old jobs destroyed. Each time, some people look to the past and assuage fears that there will be mass unemployment, while others will highlight why this time it is different. “Luddite” is now a blanket term used to describe people who dislike new technology, but its origins date back to an early 19th-century labour movement that railed against the ways that mechanized manufacturers and their unskilled labourers undermined the skilled craftsmen of the day.³¹

While earlier industrial revolutions have led to growth in employment, this document examines if that is going to be the case with the AI revolution too or will it reduce more jobs than it creates.

Through the course of this document, we seek to understand the state of AI development today, and how tasks currently being done by humans may transition to AI. In section 1, we talk about the kind of tasks that AI can do better than humans and in section 2, we list the technologies and sub-technologies involved and how they are contributing towards the disruption, while also listing the current limitations of the technologies. In section 3, we then look at India and data around employment across industries in the services sector, which we see as the sector employing the most people outside of agriculture. In section 4, we try to distinguish between AI capabilities and AI adoption and talk about the expected lag in AI adoption. In Section 5, we then develop a framework that companies choosing to develop AI solutions could use to see where AI can be effectively deployed. In sections 6 and 7, we dive deep into the Information Technology and Education sectors and show how this framework could be used to predict the transition to AI. In conclusion, we recommend an approach policymakers and those adversely impacted by AI could take to deal with job transitions, especially in the services sector.

1. What Can AI Do Better Than Humans?

At a very broad level, if we look at what AI is being used for, and consider various reports like the World Economic Forum's (WEF) Future of Jobs report²⁹, the IMF report "Gen-AI: Artificial Intelligence and the Future of Work"³⁰, etc., we can say AI sees application where there is a need to improve productivity; reduce repetitive, mundane jobs; do some things with greater accuracy; reduce the amount of risk to humans while performing the job; take data from multiple sources, access vast bodies of knowledge and make sense of extremely complex set of things.




Data processing and analysis

AI can rapidly analyse vast datasets, identify patterns that would be difficult for humans to do manually.



Repetitive tasks

Since AI does not get fatigued and it has the capacity to perform repetitive tasks with high accuracy and consistency.



Complex Calculations

AI can perform complex mathematical calculations, simulations, and optimizations much faster and more accurately than humans.



Facial and Image Recognition

AI can recognize and identify faces, objects, and patterns in images and videos with a high degree of accuracy.

**Prediction and Forecasting**

AI is good for analysing large datasets and identifying patterns. It can make accurate predictions and forecasts in areas like weather forecasting, stock market trends, and consumer behaviour.

**Personalisation and****Recommendation Systems**

AI can analyse user data and preferences to provide personalised recommendations for products, services, or content

2. What technologies are Driving AI Adoption?

In the previous section, we saw the kinds of tasks AI is good at. This section seeks to list the technologies and the sub-technologies that drive those capabilities. This section also lists some of the known limitations of these technologies to provide a better understanding of the challenges in implementing and adopting the technology.

Name of Technology	Sub-technologies	Limitations
Machine learning: Using data and algorithms to train AI to mimic human behaviour. AI learns to modify itself overtime. ³	<ul style="list-style-type: none"> a. Supervised learning b. Unsupervised learning c. Reinforcement machine learning 	<ul style="list-style-type: none"> a. Lack of transparency and interpretability, b. Bias and discrimination from data set, c. Overfitting (trained too successfully on training data), d. Underfitting (overly simplistic), e. Limited data availability, f. Needs many computational resources g. Lack of causality⁴
Natural language processing: Using machine learning to help AI in understanding and communicating in human language. ⁵	<ul style="list-style-type: none"> a. Syntactic Analysis b. Semantic Analysis 	<ul style="list-style-type: none"> a. Cannot recognise language differences, b. Training data might be biased c. Development time and resources might be many, d. Phrasing ambiguities are not recognised e. Likelihood of grammatical errors f. System contains innate biases

		<ul style="list-style-type: none"> g. Synonyms present lexical challenges h. Multilingualism cannot be recognised i. Presents false positives⁵
<p>Automated Speech Recognition: Using machine learning to convert human speech to text.⁶</p>		<ul style="list-style-type: none"> a. Long Training, b. Prolonged dictation presents challenges c. Time for templates to be developed d. Potential outages and downtime present issues e. Domain mismatch cause problems f. Vocabulary and jargon go unrecognised g. Accented speech is not picked up h. Short utterances and backchanneling is not recognised i. Punctuation remains improper⁷
<p>Computer Vision: Using machine learning to interpret visual media and detect defects with inputs.⁸</p>	<ul style="list-style-type: none"> a. Deep Learning b. Convolutional Neural Network 	<ul style="list-style-type: none"> a. Lack of specialists b. Needs regular monitoring c. Variable lighting conditions cause issues d. Perspective and scale variability cause problems e. Occlusion

		f. Contextual understanding is hard g. Lack of annotated data ⁹
Generative AI: Using training data to generate high quality text, audio and visual content. ¹⁰	a. Generative Modelling <ul style="list-style-type: none"> - Generative Adversarial Networks - Transformer Based Models 	a. Has a data dependency b. Black box nature (generate outcome or conclusion that is not transparent) c. Easily fooled d. Lack of creativity and contextual understanding e. Inability to generalise beyond training data f. High resource use g. Privacy and security are compromised h. Prone to bias ¹¹
Expert Systems: Using human reasoning through AI techniques to solve problems. ¹²	a. Forward Chaining b. Backward Chaining	a. Lack of emotions b. Common sense is compromised c. Domain specific d. Does not learn automatically (needs manual updating) e. Cannot explain logic behind decisions ¹³

Table 1 : AI Technologies

3. Distribution of Employment

Across Sectors

It may be useful to first understand how work is distributed across sectors and types of jobs before we dive deep into understanding how AI will likely make an impact. We have looked at only the services jobs in India for now since it is that sector that generates 54% of non-agricultural employment and almost 50% of total GDP.¹⁴

Since we have considered only the service sector, we don't wish to claim that the framework considered in this paper could apply to all sectors. That would be a more extensive study, and outside the scope of this document.

We have selected Information Technology (IT) and the education sector for a deep-dive, since they employ a considerable percentage of people and are also significantly different in the types of tasks done within those sectors. Together they represent a significant chunk of work done in the services sector.

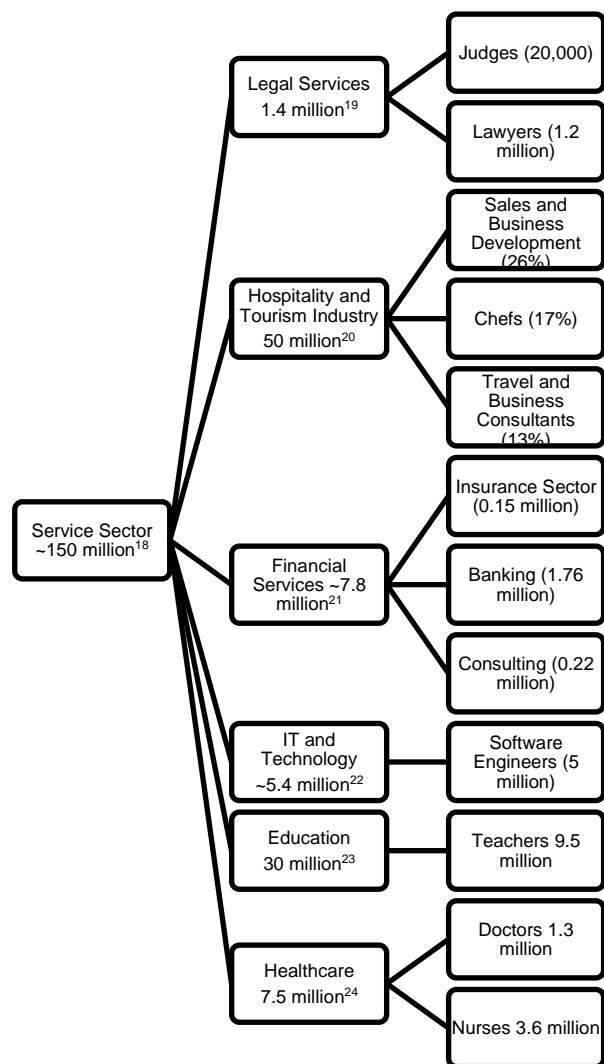


Figure 3: Distribution of jobs across various sectors

4. AI Capability vs AI Adoption

So far, we have discussed what AI can do, the technologies driving AI capabilities, and how jobs are currently distributed. This section seeks to understand how AI adoption lags behind AI capabilities and why not everything AI can do will necessarily be done by AI.

The World Economic Forum publishes a bi-annual Future of Jobs report. In 2020, the report indicated that 33% of jobs were done by machines and 67% by humans.¹⁵ It predicted that by 2025, 47% of the jobs would be done by machines.¹⁵ However, in the 2023 report (they did not publish a report in 2022), they pointed out that only 34% of the jobs were done by machines and 66% of jobs are still done by humans. They have now predicted that 42% of all tasks will be done by machines by 2027.¹⁶

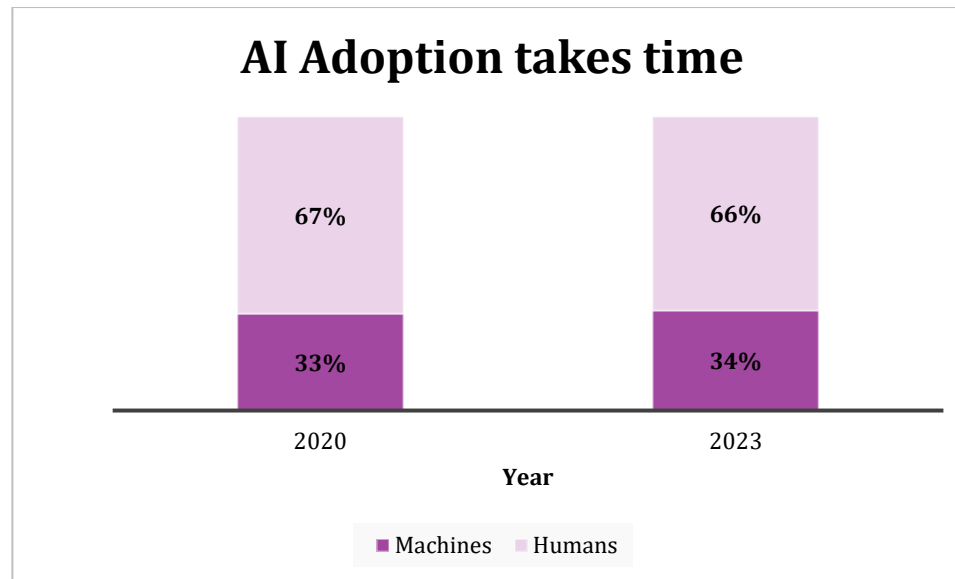


Figure 4: Task Distribution

Source: Future of Jobs report 2023 by WEF

This chart shows that while AI technology is making great leaps, the adoption of AI is not keeping pace. Even the 1% increase is not necessarily all due to AI. It could just be use of simple machines and automations. However, it is clear that AI adoption which is a subset of this increase is not progressing at the speed originally anticipated.

Bill Gates is credited with the statement, “people overestimate what they can do in one year and underestimate what can be achieved in a decade.¹⁷” This is true for AI, and here is why.

While technology capabilities may have been demonstrated, the cost of the technology is something to be considered, and its availability to those who can use it takes time. Much of the development is taking place within large corporations and is proprietary in nature.

Confidence in these technologies needs to develop, and considering the significant amount of investment that may need to be made, the changes to the processes required, and the need to understand the risks involved, the adoption of AI takes a lot more time.

There is a need to understand how the transition towards the adoption of AI may take place over time.

5. Proposed Framework for Transition to AI

As we have seen in the earlier sections, AI can do many things, but what is likely to get deployed depends largely on the areas where private enterprises see opportunities. We offer a framework to decision makers in AI development companies to prioritise the AI solutions that may get adopted. The framework takes the shape of a funnel where, at the top, there is a list of all the tasks AI can do, and at the bottom of the funnel are tasks that are likely to be done by AI.

The funnel will look like this. This framework is applicable across the different AI technologies with a specific focus on the services sector.

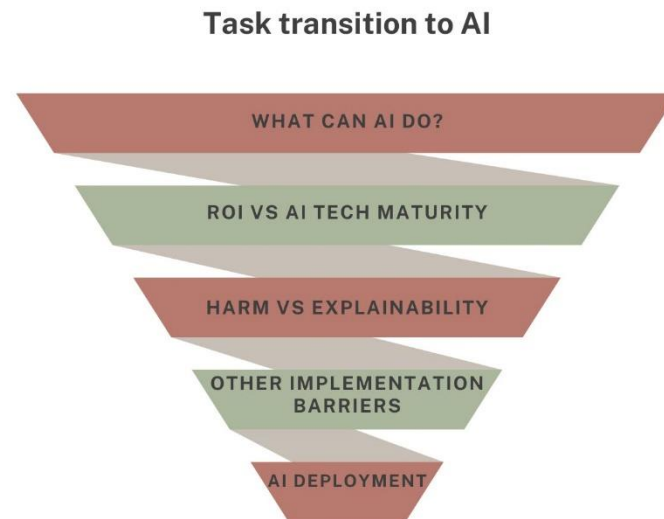


Figure 5: Funnel framework for jobs and AI

5.1 What AI can do?

As seen in the earlier sections, at the top of this funnel, there will be a list of all the tasks that AI can do from among the tasks that a job entails. This will involve looking at the technologies and sub-technologies involved in doing those tasks, and identifying the ones that AI will be able to do.

The idea is to basically think in terms of tasks, and not jobs, when it comes to this first level of the funnel.

5.2 Return on investment vs AI technology maturity

It is important for decision makers in companies to understand the technology maturity and the Return on Investment (ROI) as the next step towards moving tasks down the funnel.

List the technologies and the tasks they can perform currently and assesses the maturity of such technologies. Then assess whether AI can do things better than humans. It is not enough if it can do them. It needs to do them cheaper, faster, more accurately or more creatively.

It is then important to understand the investments required to deploy the AI and whether it will be more efficient in the run phase. This will be used to calculate the ROI.

The chart below shows how ROI and AI maturity can help shortlist tasks for AI deployment initially.

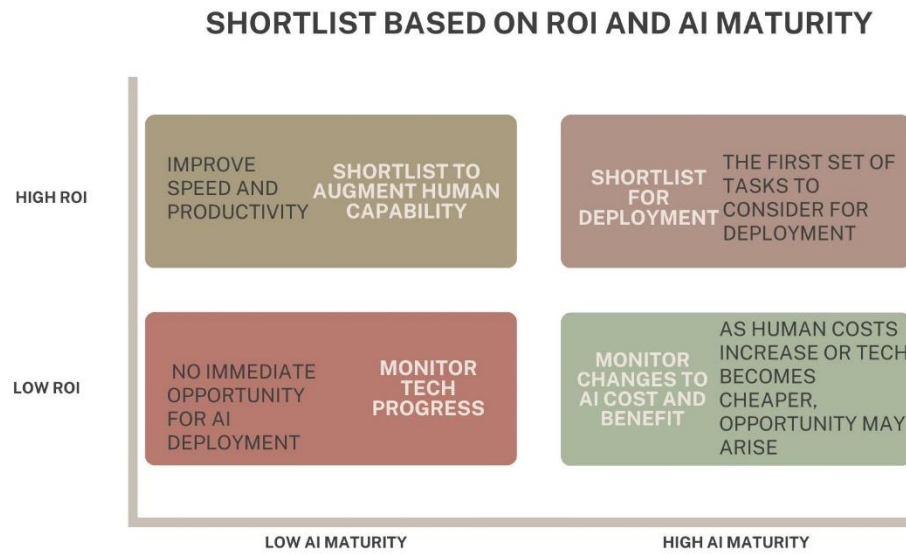


Figure 6: ROI and AI Maturity Chart

Take the example of call centre agents who respond to customer queries from healthcare insurance policyholders. The tasks they do are as follows:

1. Say hello.
2. Ask for a policy number and other customer information.
3. Listen to the query regarding coverage or status of claim.
4. Check with the policy document and confirm whether a particular treatment is covered under the policy or check with the workflow system within the insurance company and inform and update the customer on the status of their claim.

5. Listen to the customer and assess if they are happy or unhappy with the response.
6. Pass on the call to a senior person on their team if the customer is unhappy.

The AI technology that can be used to handle this is fairly straightforward and quite mature. Generative AI tools can be trained on a small language model to answer questions and even ask follow-up questions. The cost of such AI is also not significant compared to hiring humans to do this task, and therefore it makes sense to shortlist this application for AI deployment.

- High ROI, low technology maturity: Online healthcare
 - o There is a good use here because healthcare is underfunded, but the AI technology for online healthcare has not advanced enough for it to be widely adopted; it still needs more research and funding.
- Low ROI, high technology maturity: High end culinary arts
 - o AI can contribute to art and design but the cost of developing AI that can replicate human emotion and depth will be too expensive and the returns would be low compared to the investment required. Hence while the technology is capable of generating pretty impressive art and design, it is unlikely to replace human artists immediately.
- Low ROI, low technology maturity: Autonomous robot home chef

- The technology will be difficult to adopt and people probably prefer manually prepared kitchen food and it's also easier, so it will not be adopted. It also involves other logistical issues around acquiring material required for the process, and the element of taste variance that cannot be accounted for always. It would be safe to assume that the low technology maturity with respect to the entire process may inhibit its widespread adoption, at least for now. Moreover, unless it is completely autonomous throughout the entire process, it doesn't provide attractive ROI when it comes to time and effort saved.
- High ROI, high technology maturity: Fraud detection systems in banking and financial services
 - There is a considerably high technological maturity achieved in these systems already through years of development in machine learning algorithms, pattern recognition and real-time data processing. The ROI in terms of preventing fraudulent transactions, reducing false positives, increasing operational efficiency, etc. is also substantially high. Consequently, these technologies are a standard component of the security infrastructure in the banking and financial services industry.

Once the organisation shortlists a bunch of tasks using this set of filters, it is time to go to the next step.

5.3 Harm vs explainability

The next level of short listing involves estimating the harm that errors in doing the task can cause and pit them against the ability to explain/verify the outcome. Something that is easy to explain/verify and that is likely to cause minimal harm, even in case of errors, is more likely to be accepted by society. Such tasks will fall into the first category of tasks that AI may take over. The decision maker will look here for acceptability of the solution to users.

For a task to be considered acceptable to users of AI, there are two significant hurdles to overcome. What is the probability that the AI will make an error, and in case of the error, what harm will be caused? The product of the probability of the error and the impact of the error is expected harm. It is obvious that tasks, which if done by AI have high expected harm, are unlikely to be acceptable to users unless there is a human in the loop who can be relied upon to catch the error and prevent the harm. And vice versa is also true. Tasks that have low expected harm are more acceptable to the users.

However, there is another filter to look at as well. It may not be enough if the AI does not cause harm but there should be a way by which we know the AI is not wrong. This is what we mean when we say apply the filter of explainability or verifiability. Black box AI solutions are less likely to be acceptable compared to AI solutions that can be explained or verified.

The chart below shows how these two filters of explainability/ verifiability and expected harm will be applied to come up with AI that can be deployed at least from the user's point of view.



Figure 7: Explainability/verifiability AI chart

5.4 Other implementation barriers

Once these basic filters are crossed, there is a list of tasks AI is likely to end up doing but there are some more barriers to implementation in the real world.

While AI, under these circumstances, may appear to be useful and acceptable to the users, there are those whose livelihoods may be impacted by such AI deployment. Depending on their bargaining power, governments may impose restrictions on AI deployment, at least, in the short run.

AI solutions will have to face the challenge imposed by labour unions that may oppose the loss of jobs caused by AI. The bargaining power of these unions can be quite considerable. The cost of compensating the people adversely affected may change the ROI calculations significantly. There is also the political fallout of large-scale unemployment to evaluate and deal with. Different industries and countries may be impacted by this in different ways. This is likely to lead to lags in AI deployment in some industries and countries. The Union Transport Minister Nitin Gadkari is credited with saying, “As long as I am there, I will never allow driverless cars in India. I said in the US itself that at any cost, I won't allow such cars in India. Otherwise, 80 lakh drivers will lose jobs²⁸.”

Then there are the industry bodies with entrenched players finding it difficult to compete with new companies that come in and disrupt the industry with AI based solutions. These industry bodies may have relationships with the political establishment and use industrial policy as a weapon to slow AI deployment. The automobile industry in the US has been seeking guidelines from Congress on issues on safety and liability and the intention here seems to be to slow down the adoption of autonomous vehicles before the entrenched players can get up to speed.³²

They may also join forces with labour unions and question the need for AI. They may have the resources to rally the public against AI.

Some countries may see AI technology as a weapon and try to prevent other countries from developing similar capabilities. Technology available in one country could get restricted from being shared elsewhere under the garb of national interest and national security. A case in point is the US banning exports of certain AI chips to China.

5.5 AI deployment

The tasks, after flowing through the above levels in the funnel, will give us an idea about the scope of adoption of AI in those particular tasks and jobs, and the nature of their deployment. Based on where they get placed in the

frameworks at levels 2 and 3, and how they fare against other implementation barriers, this final level will be the AI deployment level.

6. Deepdive into Information Technology Sector

This section takes the framework recommended in the previous section and applies it to the Information Technology sector as it exists in India. Here we examine how the sector may be impacted by deployment of AI. While this analysis applies in a generic fashion globally to the sector and its roles and technologies involved, we have tried to apply the lens to India specifically whenever possible.

This sector employs approximately 5.4 million people in India, out of which approximately 5 million are software engineers alone.

Let's look at some of the tasks that we might find generally in this sector through the funnel explained above for the potential impact of AI on jobs in the information technology sector in India:

Role	Task	Technology	AI Tech Maturity	ROI	Explainability	Harm	Replace/Augment/For Later
Software Engineers	Code generation, bug detection, automated testing	NLP, code analysis tools	High	High	Medium	Over-reliance on AI-generated code, potential security vulnerabilities	Augment
Software Engineers	System architecture design and optimisation	AI-powered design tools	Medium	Medium	Low	Inefficient solutions if not properly reviewed	Augment
Production Support	System monitoring, alert generation, basic ticket resolution	AI monitoring tools, automated problem-solving	High	High	Medium	Missed critical issues if not properly trained	Replace (for routine tasks)
UX designers	UI mockups, design improvements, basic UI elements	AI design tools	Medium	Medium	Medium	Lack of human touch and cultural sensitivity	Augment

Product Managers	Analyse user data, market trends, competitor information	AI data analysis tools	Medium	Medium	Low	Over-reliance on AI-driven insights, lack of innovation	Augment
Product Managers	Product health analysis	AI analytics tools	Medium	Medium	Medium	Misalignment with user needs	Augment

Table 2: Impact of AI on jobs – Funnel

Note: In most cases, AI is seen as augmenting human roles rather than fully replacing them, especially in areas requiring creativity, complex decision-making, or human intuition. The "For later" option wasn't explicitly used as these AI applications are currently viable, though their implementation may be gradual.

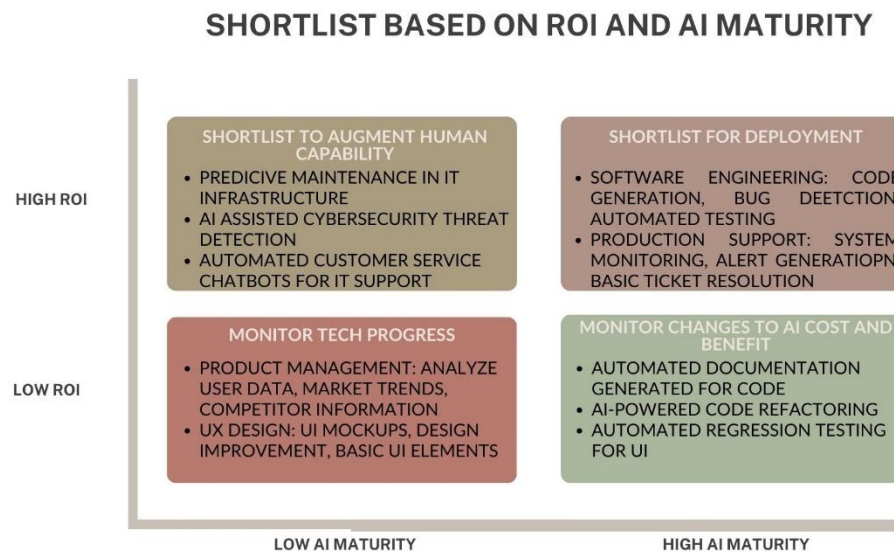


Figure 8: Framework application for IT sector

AI's impact on IT roles is significant and multifaceted. In software engineering and production support, AI can automate routine tasks, enhance code generation, and streamline system monitoring, offering high ROI due to the relative maturity of these technologies. For example, there are already tools like GitHub Copilot for enhanced code generation. For UX design and product management, AI can provide valuable insights and assist in basic tasks, though with potentially lower ROI due to the need for human creativity and oversight in these jobs. However, the implementation of AI across these areas comes with risks, including over-reliance on AI-generated

solutions, potential security vulnerabilities, and the loss of human intuition in design and product decisions.

If we look at India specifically, the deployment of AI in IT roles faces a few barriers. While traditional labour unions are less prevalent in the IT sector, professional associations may voice concerns about job displacement. Industry lobbies might advocate for gradual AI adoption to allow for workforce reskilling. Mid-level managers may resist AI implementation due to fears of obsolescence. These misgivings can be handled by providing upskilling and reskilling opportunities, because most of the job changes will entail this. Also, IT is perhaps one of those industries which are very adaptable to change; the inertia against change in this sector is not a lot, as showcased by India's growing prowess in the software industry.

Moreover, despite these challenges, the ease of AI deployment varies across roles, with software engineering and production support offering relatively straightforward integration opportunities, while UX design and product management may require more customisation and human oversight in AI tool implementation.

The impact of AI on India's IT sector is, therefore, likely to be significant but uneven across different roles. Software engineering and production support roles are at higher risk of automation, potentially affecting a large portion of

the 5 million software engineers. However, this doesn't necessarily mean job losses, but rather a shift in required skills.

UX designers and product managers may see their roles evolve rather than be replaced, with AI augmenting their capabilities rather than replacing them entirely. The key challenge for India's IT sector will be to manage this transition, focusing on reskilling and upskilling the workforce to work alongside AI systems.

Given the size of India's IT workforce (~5.4 million), even a small percentage of job displacement could affect hundreds of thousands of workers. However, the implementation barriers and the need for human oversight in AI systems may slow this transition, providing time for adaptation.

The explainability of AI used would be higher at lower levels of implementation, like at the software developer's level where the high-level design or architecture is not worked upon directly. For tasks requiring a higher level of understanding of the system, and its integrations and interactions with outside systems, the explainability might be a little difficult because it might not be a simple, straightforward algorithm to implement. This will then require a collaborative effort between humans and AI. It requires intuition, an ability to work at the abstract level, and to take a holistic view of the system under consideration.

Misalignment with user needs could become a potential avenue for harm. Therefore, at a higher level of abstraction, an augmentation of the AI capabilities with human oversight is needed. This will also be needed at the architecture and design level to prevent security issues, which are not just related to secure coding practices in isolation, but also to system integration. Automated code review tools help at the lower levels; but at higher levels of abstraction, we need more experience which AI tools may not always be trained adequately on.

The IT industry in India will likely need to focus on developing higher-value skills that complement AI capabilities, such as AI system design, ethical AI implementation, and complex problem-solving that AI cannot yet handle. This shift could potentially maintain India's competitive edge in the global IT market while transitioning to a more AI-augmented workforce.

Let's try to gauge the impact of AI on jobs in the IT sector.

Please note that these are approximations and actual impact may vary based on numerous factors including the pace of AI adoption, regulatory changes, and economic conditions. These are very dynamic, influencing vectors, and their impact can only be approximated. What could be done to minimise their impact is to reduce the uncertainty and variability in these factors. With the changing nature of AI, the world is always going to be in flux when it comes to policies surrounding AI and its applications; hence, this is indeed

easier said than done. Keeping all stakeholders involved while making decisions regarding regulations about AI will go a long way in reducing uncertainty and variability.

Let's break down the potential impact based on the general industry trends with regard to experience levels:

1. Junior developers (1-4 years of experience)
 1. This group might be most vulnerable to AI impact, as many of their tasks involve more routine coding that AI could potentially handle.
 2. Potential impact: Moderate to High
2. Mid-level developers (5-9 years of experience)
 1. This group might see a mixed impact. While AI could automate some of their tasks, their experience would still be valuable for complex problem-solving and overseeing AI-generated code.
 2. Potential impact: Moderate
2. Senior developers (10-14 years of experience)
 1. This group might be less impacted directly by AI, as their roles often involve high-level design, architecture, and complex problem-solving that AI is not yet adept at.
 2. Potential impact: Low to Moderate
2. Other experience levels (assumed to be very junior or very senior)

1. The impact on this group would vary greatly. Very junior developers might face challenges similar to the 1-4 year group, while very senior developers might be less directly impacted.
2. Potential impact: Varies widely

Key Points:

1. The impact is likely to be most significant for junior developers, potentially affecting their entry into the field or necessitating rapid upskilling.
2. Mid-level developers may need to adapt their skills to focus more on overseeing and refining AI-generated code, as well as tackling more complex problems.
3. Senior developers may be less directly impacted but will likely need to develop skills in AI integration and management.
4. The overall impact suggests a significant transformation of the software engineering field, rather than wholesale job losses.

It's crucial to note that this analysis assumes AI capabilities continue to advance and are widely adopted. The actual impact could be moderated by factors such as:

- The pace of AI adoption in Indian companies
- The emergence of new roles related to AI oversight and integration
- The ability of the workforce to upskill and adapt to working alongside AI tools

- Potential increases in overall demand for software development services

It's crucial to note that "jobs affected" doesn't necessarily mean job losses. Many of these roles will likely evolve to work alongside AI, requiring reskilling and upskilling. Additionally, new roles may emerge to manage and oversee AI systems, potentially offsetting some of the impact.

The IT industry in India will likely need to focus on transitioning a significant portion of its workforce to more AI-complementary roles over the coming years to maintain its competitive edge and adapt to these technological changes.

7. Deepdive into Education Sector

This section examines the impact AI may have on jobs in the Education Sector in India similar to the earlier section where we examined the impact of AI on the Information Technology Sector in India.

Total number of employees: ~30 million

Some important jobs in the education sector, along with the list of their tasks and job requirements, are:

Role	Task	Technology	AI Tech Maturity	ROI	Explainability	Harm	Replace/Augment/For Later
Teacher	Imparting knowledge, providing resources, curriculum creation, setting goals and tracking progress, grading, designing tests, administrative tasks, parent communication, planning school events	AI-powered learning platforms, digital grading tools	Medium	Medium	Low	Potential loss of personal touch	Augment
Registrar	Student registration, data entry, paperwork processing,	Student Information Systems (SIS),	High	High	High	Data security risks, potential for errors in	Replace

	financial management, appointment management, record maintenance, record organisation, complaint management	Database Management Systems				automated systems	
School Counsellor	Arranging college visits, post-school planning, emotional support, career advising, communication with teachers/parents, college application assistance, conducting assessments	AI-powered career matching tools, online counselling platforms	Medium	Low	High	Potential misinterpretation of AI-generated advice, lack of human empathy	Augment

Admissions Officer	Creating acceptance criteria, filtering student profiles, student advising, communicating decisions	AI-powered applicant screening tools, Customer Relationship Management (CRM) systems	Medium	High	High	Potential bias in AI-driven decision making, lack of holistic evaluation	Augment
--------------------	---	--	--------	------	------	--	---------

Table 3: List of tasks in Education Sector

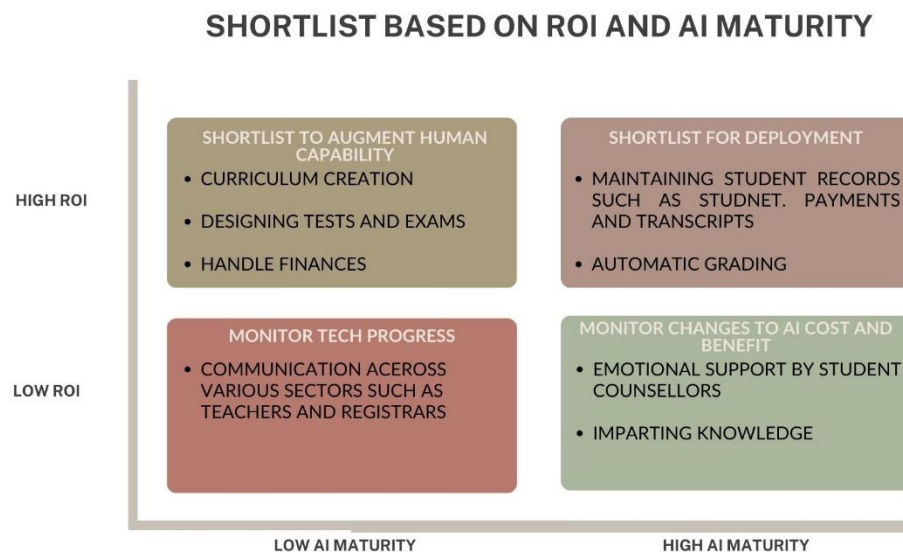


Figure 9: Framework application for Education sector

It is easy to see how certain tasks fit quite seamlessly into the criteria to be replaced by AI. Jobs that require excessive manual labour, or are not associated with sentiment and emotion, and for which technology exists are likely to be the first to be replaced by AI.

Naturally a big area where AI has found its way is the job of the registrar. Much of the jobs that were earlier done manually by registrars have now been overtaken by AI. Enrolment, scheduling and transcript management are not only being done more efficiently, but more accurately by these new systems²⁵. Moreover, the inclusion of AI into the system has improved student compliance, by eliminating bias from the system. After passing the job of the registrar through the earlier proposed frameworks, it seems that the task has a high ROI and low stakes since registrars do not do much more than clerical work. Moreover, the explainability of such a system is high, since registrars mainly engage in procedural tasks. Although it is unlikely that registrars will be completely replaced by AI in the near future, this is one potential area where there is a scope for automation.

Tools such as predictive software have already found their way into many education systems. Top-ranking colleges have now incorporated these mechanisms as a part of their admissions departments to analyse college applications. Roughly 60% of college admissions officers use AI to analyse essays, and 50% use it to conduct college interviews²⁴. Another 61% of admissions officers use AI to communicate with their applicants. High-

ranking universities internationally, such as the University of Pennsylvania in the USA, have already incorporated AI technology into their systems. They used AI to analyse vast datasets to come up with seven desirable qualities they seek in applicants. Analysing over 300,000 essays, their AI system named RoBERTa was able to accurately identify potential admits, as verified by 36 admissions officers.²⁶

Along with analysing student applications to universities, AI has now found its way into the jobs of college counsellors as well. Students are now using online systems to guide their college applications. Online platforms such as Niche and CollegeVine use AI technology to determine a student's chance of being admitted into a particular university, based on analysis of vast datasets of other students who have been admitted to the same university. These platforms even have the capability to compare a student's own profile to the other profiles of students who have been admitted based on extracurricular activities and grades.

However, in both these cases, the stakes for jobs to be completely replaced by AI are high, raising big ethical questions. Highly selective institutions, with admission rates less than ten percent internationally, have maintained that their college admissions require nuanced insight that AI cannot provide ²⁷. Even though the efficiency provided by these systems has streamlined the process, a high degree of oversight is still needed to make these decisions. In the case of counsellors, while college admissions serve a big part of their role,

they do serve other tasks associated with sentiment and emotions. Students use these individuals as moral support and a means for advice that AI cannot provide. The personal touch that counsellors provide cannot be done by any machine.

Conclusion

The tasks AI is likely to replace will be determined by the progress of the technology but also by the return on investment, potential for harm, explainability and verifiability. It will also have to deal with the bargaining power of those whose status quo is being disturbed. It is likely to however eventually prevail.

Policymakers in India could look to the EU for some guidance here. We need areas of AI that may be unregulated if the potential for harm is negligible, while those with significant potential harm should, at the very least, be explainable or verifiable.

Governments should not try to block AI deployment only because such deployment could lead to job losses. A local government in prehistoric times, if it existed, might have, using that logic, prevented the use of a wheel because it now makes those who carry heavy loads lose their jobs and mankind would have perished.

The education system needs to be revamped to ensure one is prepared for a lifetime of learning. Teaching people how to learn will be the most important task of a school system.

Upskilling and reskilling are likely to be the order of the day. Huge amounts of resources may be required to upskill and reskill people, but using tools like career impact bonds will allow these to be driven by the market. There is a place for philanthropy and government subsidies, but they could be done by subscribing to equity tranches of career impact bonds rather than by governments and even individuals trying to determine who should be trained on what.

This could also act as safety nets that provide for loans to people between jobs to acquire new skills and even feed themselves while getting themselves more relevant in the jobs market.

AI is likely to displace jobs at a large scale. It will, however, take time. There is a need for prudence, patience and acceptance of the changes that are coming. Continuous learning, upskilling and reskilling will be the normal way to live.

Career Impact bonds are financial instruments and the term was coined by a US based company called Social Finance.

Here seekers of upskilling and reskilling are offered an at-risk loan which is repayable only if they get a job post the training with an income above a particular threshold.

These loans are pooled and securitized. The bonds are divided into multiple tranches and the cash flow from the repayment of the loans and the interest on the loans are directed into the different tranches based on seniority. The senior tranches get the money first and are less risky than the bottom equity tranches.

This ensures the interests of the skill-seeker, the training institute and the financial investors are all aligned and funding will go towards training programs and institutes that are likely to return a positive outcome.

References

1. Giattino, Charlie, Mathieu, Edouard, Samborska, Veronika, and Roser, Max. "Artificial Intelligence." Our World in Data, August 16, 2023. <https://ourworldindata.org/artificial-intelligence>.
2. Karabegović, Isak & Ćar, Mirha & Šestić, Munira & Stupar, Savo & Karabegović, Edina & Husak, Ermin & Mahmić, Mehmed & Isić, Safet & Vojic, Samir & Banjanovic-Mehmedovic, Lejla & Buljan, Stipo & Lemeš, Samir & Zaimović-Uzunović, Nermina. (2022). Industry of Bosnia and Herzegovina within Industry 4.0. 10.5644/PI2022.202.20.
3. n.d. What is machine learning (ML)? Accessed July 16, 2024. <https://www.ibm.com/topics/machine-learning>.
4. Tabrez, Sohail. Limitations of Machine Learning. Accessed July 14, 2024. <https://www.tutorialspoint.com/7-major-limitations-of-machine-learning>.
5. Holdsworth, Jim. What is NLP (natural language processing)? July 16, 2024. <https://www.ibm.com/topics/natural-language-processing>.
6. Major Challenges of Natural Language Processing. March 12, 2024. Accessed July 14, 2024. <https://www.geeksforgeeks.org/major-challenges-of-natural-language-processing/>.
7. LabCE, Limitations of Speech Recognition Technology. Accessed July 14, 2024. https://www.labce.com/spg3231359_limitations_of_speech_recognition_technology.aspx.
8. Chooch. Common Problems with Computer Vision and their Solutions. Accessed July 14, 2024. <https://www.linkedin.com/pulse/5-common-problems-computer-vision-solutions-chooch-efave/>.
9. Martineau, Kim. What is generative AI? April 20, 2024. Accessed July 14, 2024. <https://research.ibm.com/blog/what-is-generative-AI>.
10. Fauscette, Michael. Understanding the Limitations and Challenges of Generative AI. January 3, 2024. Accessed July 14, 2024. <https://em36otech.com/tech-article/understanding-limitations-and-challenges-generative-ai>.

11. Zwass, Vladimir. Expert System. Accessed July 16, 2024. <https://www.britannica.com/technology/expert-system>.
12. Bhat, Samudyata. What are Expert Systems in Artificial Intelligence? April 30, 2024. Accessed July 14, 2024. <https://www.mygreatlearning.com/blog/expert-systems-in-artificial-intelligence/#limitations-of-expert-systems>.
13. Government, PIB. Ministry of Finance. January 31, 2024. Accessed July 16, 2024. <https://pib.gov.in/PressReleasePage.aspx?PRID=1793804>.
14. Schwab, Klaus, and Saadia Zahidi. "The Future of Jobs Report 2020." World Economic Forum. October 20, 2020. Accessed July 2024, 2024. <https://www.weforum.org/publications/the-future-of-jobs-report-2020/>.
15. Zahidi, Saadia. "Future of Jobs Report 2023." World Economic Forum. Accessed July 16, 2024. https://www3.weforum.org/docs/WEF_Future_of_Jobs_2023.pdf.
16. Gates, Bill. "Goodreads". Accessed July 16, 2024. <https://www.goodreads.com/quotes/302999-most-people-overestimate-what-they-can-do-in-one-year>.
17. Statista. Number of people employed across the service sector in India from financial year 2017 to 2023, by service type. Accessed July 16, 2024. <https://www.statista.com/statistics/1284079/india-service-sector-employment-by-service-type/>.
18. How many Advocates/Lawyers in India. October 17, 2024. Accessed July 16, 2024. <https://advocatedelhi.wordpress.com/2022/10/17/how-many-advocates-lawyers-in-india/>.
19. Statista. Number of direct and indirect jobs in the travel and tourism sector across India from financial year 2014 to 2023, with projection for 2024. Accessed June 16, 2024. <https://www.statista.com/statistics/1094224/india-employment-in-travel-and-tourism-sector/>.
20. Banking industry in India - statistics & facts. Accessed June 16, 2024. <https://www.statista.com/topics/5362/banking-industry-in-india/#editorsPicks>.
21. Employment of the IT-BPM industry in India from financial year 2009 to 2022, with an estimate for 2023. Accessed June 16, 2024.

- <https://www.statista.com/statistics/320729/india-it-industry-direct-indirect-employment/>.
22. Times of India. Total employment in 9 select sectors including Education, rises to 3.10 cr in July-Sep 2021. January 10, 2024. Accessed June 16, 2024. <https://timesofindia.indiatimes.com/education/news/total-employment-in-9-select-sectors-including-education-rises-to-3-10-cr-in-july-sep-2021/articleshow/88806118.cms>.
 23. Hindu, The. Demand for healthcare workers in India grew 22% during 2021-23: Study. April 27, 2024. Accessed June 2024, 16. <https://www.thehindubusinessline.com/news/demand-for-healthcare-workers-in-india-grew-22-during-2021-23-study/article66784701.ece>.
 24. Is AI affecting college admissions? | best colleges | U.S. news. Accessed July 22, 2024. <https://www.usnews.com/education/best-colleges/articles/is-ai-affecting-college-admissions>.
 25. oneorigin_admin. The Evolution of College Registrars: Past, Present, and Future. January 9, 2024. <https://www.heysia.ai/college-registrars-evolution/>.
 26. Hawley, Mitchell. Universities Are Using AI in the Admissions Process: Here's What You Need to Know. Verge AI - Unleash the Power of AI, April 23, 2024. <https://verge-ai.com/blog/how-ai-is-being-used-for-university-admissions/>.
 27. ibid
 28. Singh, Shubham. Exclusive: 'Driverless Cars Will Never Come to India,' Says Nitin Gadkari at Zero Mile Samvad. Business Today, December 16, 2023. <https://www.businesstoday.in/auto/story/exclusive-driverless-cars-will-never-come-to-india-nitin-gadkari-at-zero-mile-samvad-409765-2023-12-16>.
 29. World Economic Forum. The Future of Jobs Report 2023. Accessed August 6, 2024. <https://www.weforum.org/publications/the-future-of-jobs-report-2023/>.
 30. Manyika, James and Sneider, Kevin. AI, Automation, and the Future of Work: Ten Things to Solve For. McKinsey & Company, June 1, 2018. <https://www.mckinsey.com/featured-insights/future-of-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for>.

31. Andrews, Evan. Who Were the Luddites? August 21, 2023.
<https://www.history.com/news/who-were-the-luddites>
32. Krisher, Tom and Khalil, Ashraf. Congress urged to revive long-stalled debate about regulating self-driving vehicles. July 27, 2023.
<https://apnews.com/article/autonomous-vehicles-self-driving-cars-avs-f1cef9b9595a4bd87287670eaa1694d6>



The Takshashila Institution is an independent centre for research and education in public policy. It is a non-partisan, non-profit organisation that advocates the values of freedom, openness, tolerance, pluralism, and responsible citizenship. It seeks to transform India through better public policies, bridging the governance gap by developing better public servants, civil society leaders, professionals, and informed citizens.

Takshashila creates change by connecting good people, to good ideas and good networks. It produces independent policy research in a number of areas of governance, it grooms civic leaders through its online education programmes and engages in public discourse through its publications and digital media.