AMERICANS *for* RESPONSIBLE INNOVATION

WORKFORCE & AI REPORT
April 7, 2025

The Technologist-Economist Disconnect: Assessing AI Labor Disruption

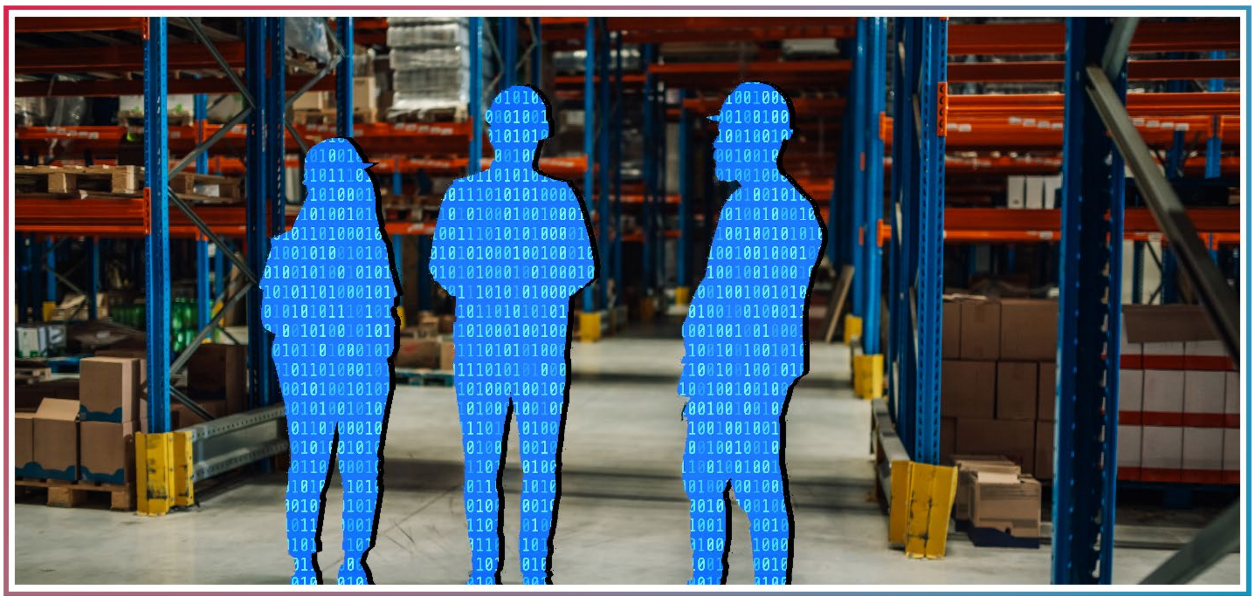

David Robusto, Policy Analyst

Since ChatGPT's debut in 2022, AI has evolved from novelty to industry necessity, accelerated by the rise of agentic systems. Simultaneously, a troubling disconnect has emerged in conversations about the workforce impacts of the upcoming change. Most economists project manageable labor market changes based on historical patterns, but leading AI developers are forecasting that transformative capabilities – such as systems that could outperform humans across virtually all tasks – will emerge as soon as 2027.

This report posits that this gap can partially be explained by a shortcoming of applying the General Purpose Technology framework to AI, and presents the concept of a "Dynamic General Purpose Technology", which better elicits AI's novelty. See Table 1 for a brief description of the three components included in this expanded framework, including Pace, Pervasiveness, and Plasticity.

Table 1

Framework for Dynamic General Purpose Technology
Pace: Describes how quickly a technology develops and is adopted.
Pervasiveness: Looks at how many industries a technology affects, its number of use cases, and how diverse the types of tasks it can handle are.
Plasticity: Evaluates how adaptable a technology is after its initial introduction and how much effort is required to apply it to new tasks, domains, and use cases.

This report also evaluates current AI workforce policy proposals and highlights four more feasible approaches that have the potential to mitigate some of the labor impacts of AI:

- **Upskilling/Reskilling**: Programs, crafted in conjunction with industry or labor unions, that help workers transition from affected jobs to new positions in automation-resilient industries.
- **Modernized Unemployment Insurance (UI)**: Adjustments to UI eligibility, benefit amount, or duration could soften automation's impact, help workers transition, and disincentivize mass displacement.
- **Vocational Schooling**: Educational pathways for automation-resilient vocational work such as plumbing and construction. Curricula developed as a partnership between government and industry.
- **Lifelong Learning Accounts**: Tax-advantaged savings accounts for education and reskilling programs. Similar to 529 education savings plans or HSA plans.

Incremental proposals like these are an important first step, but historically, major economic transitions have required bold policy interventions such as the GI Bill or various New Deal programs. Bridging the gap between technologists' grand predictions and economists' gradualist assumptions requires multiple bold policy innovations commensurate with the potential scale of disruption.

Background

"The factory of the future will have only two employees, a man and a dog. The man will be there to feed the dog. The dog will be there to keep the man from touching the equipment." –Warren Bennis

This old quip about automation has caused nervous laughter for years, but for those listening to AI technologists today – particularly as agentic AI becomes more prevalent – the empty factory no longer sounds so hyperbolic. Earlier this year, Anthropic CEO Dario Amodei told the Wall Street Journal that he believes AI will surpass humans on virtually all tasks as soon as 2027. Google DeepMind's CEO, Demis Hassabis, has similarly said that we are only a handful of years from Artificial General Intelligence (AGI). Predictions like these are not new, and may be off the mark, but companies are starting to test ambitious products that could have labor impacts. Just this month, *The Information* reported that OpenAI told investors they plan to sell advanced coding and "PhD-level" research agents – the time is right to prepare.

Despite this, ask economists about AI's impact on jobs and many will shrug, pointing to centuries of technological advancement that eventually created more employment and prosperity than they displaced. However, the very architects of this technology predict unprecedented transformation within just the next few years. If their vision comes to pass, human cognitive labor could become redundant, fundamentally reshaping the workforce, economy, and society as we know it. Even a moderate probability of this scenario warrants taking these technologists seriously, which means exploring why their perspective differs from those who point to historical technologies, as well as developing robust policy responses before disruption occurs.

Why AI Might be Different

Establishing Dynamic General Purpose Technologies

Economists have increasingly begun to view generative AI through the lens of a General Purpose Technology (GPT)[1] – a framework that has helped explain diverse transformative inventions from steam power to electricity to the internet. While this label clearly does apply to AI – there is no doubting the technology's rate of development and broad impact across sectors – it may not fully capture what makes AI so economically impactful. I propose extending the GPT framework to consider *Dynamic* General Purpose Technologies (DGPTs) by adding a novel factor, plasticity. The three components of a Dynamic General Purpose Technology, explained through an evaluation of AI, are as follows:

Pace – Ever since the release of ChatGPT, the development and adoption of generative AI has been moving at breakneck speed. Since 2022, frontier models have become multimodal (text, images, audio, video), improved their reasoning capabilities, gained external tool integrations, started to incorporate agentic capabilities, and more. On the adoption side, ChatGPT is the second fastest app ever to reach 100M users, reaching the milestone in just two months. The rapid uptake has been similar in industry – in 2023, a McKinsey survey indicated that 33% of businesses were already using the technology regularly in at least one function. As of July 2024, an updated survey puts that number at over 70%.

Compare this sharp uptake to previous technological revolutions that unfolded over decades or generations – the Industrial Revolution took nearly a century to fully materialize. One way to internalize generative AI's rate of progress is by calculating how long it took to go from its Main Breakthrough to Widespread Adoption (MBWA). What constitutes the "main breakthrough" for any of these technologies can be seriously contested, but is loosely defined by asking "what moment most enabled this technology to reach mass scale?" Similarly, "widespread adoption" for any of these technologies is ambiguous and often

[1] Scholars define a General Purpose Technology as a technology (to paraphrase Lipsey et al.) that has significant room to evolve from its inception, is very widely used, has a massive number of use cases, and has many cascading impacts or externalities.

based on limited data, but is estimated through historical rates of diffusion and adoption.

The table below compares the MBWA to five transformative historical technologies, with this new tech being decidedly the fastest. See the Appendix for a full breakdown.

Table 2

Technology	Main Breakthrough	MBWA
Steam Engine	James Watt developed his improved steam engine, patented in 1769.	~70 years
Spinning Jenny	The exact origin point of the Jenny is debated but it was created by James Hargreaves in the mid 1760s.	~20 years
Steel	Henry Bessemer created his eponymous process in 1856 which significantly simplified and cheapened steel production.	~15 years
Electric Motor/Electricity	The first practical DC motor was introduced by Frank J. Sprague in 1886.	~30 years
The Internet	In 1983, ARPANET adopted the TCP/IP protocol.	~12 years
Generative AI	In 2017, Google researchers released the paper "Attention is All You Need," which introduced the Transformer architecture.	~5 years

Pervasiveness – Like other GPTs, AI is impacting widespread fields from medicine, to manufacturing, to art. What is different is that generative AI is already demonstrating transformative potential in both physical labor and cognitive work, where most technological revolutions of the past have been either-or. The steam engine transformed manual labor, but had little effect on cognitive tasks, while the internet revolutionized cognitive work, but has had less impact on physical tasks. A full breakdown in the Appendix helps paint the historical picture.

AI breaks this pattern through its potential for simultaneous impact on both the physical world (e.g., robotic autonomy) and complex cognitive functions (e.g., creative writing). Moravec's paradox – the observation that computers excel at cognitive tasks like reasoning but require much more computation for physical capabilities like motor skills – has resulted in many physical tasks being deemed "automation-resilient." But that does not mean such tasks are immune over time. AI's unique historical position on this dimension – especially when considered in conjunction with the final factor – shows that a world where these industry insiders are correct is a very uncertain place for human labor.

Plasticity – The final factor is a novel addition to this framework and the general economic conversation – plasticity, which captures how adaptable a technology is after its initial introduction. Fixed-function technologies of the past, like the steam engine and the spinning jenny, could only perform the specific tasks they were engineered for; their new applications expanded through complementary innovations. While certain GPTs are not fully static – such as the internet, which has been re-architected to handle things like real-time applications – their adaptations are iterative and require significant human effort.

Modern AI, in contrast, can be more fundamentally reconfigured to handle novel modalities like images, and capabilities like reasoning via chain-of-thought. These systems can also be simply retrained or "fine-tuned" to handle new use cases. Theorized AGI systems could take this even further, able to independently extrapolate to novel tasks and potentially exhibit behavior like recursive self-improvement (p. 63).

This has profound implications for labor markets: previous technological shifts created new job categories that held stable for decades, such as railroad workers or electrical engineers. AI can potentially learn to perform the very tasks that its existence creates, which threatens to disrupt employment in continuous waves to a degree never seen before.

For an example of how accounting for plasticity would shift current predictions, consider McKinsey's 2023 report[2] estimating the impact of AI automation on labor demand. McKinsey predicts that automation, including AI, will have a massive impact on the U.S. workforce in the near future, taking over tasks that account for roughly 30 percent of the hours worked in the economy by 2030. But, the report notes, job losses due to automation are likely to be offset by the creation of new occupations and "work activities," with McKinsey predicting that the occupational categories most susceptible to generative AI automation will see net growth in labor demand through 2030.

This makes sense if you assume a strong reinstatement effect[3] – as has been the case with many previous waves of automation – but high plasticity would mitigate or even eliminate this. If many of these newly created "work activities" could also be automated, McKinsey's projected numbers would need to be considerably lower, although how much lower varies based on the rate of technical progress and adoption.[4]

Transformative Technology Requires Bold Policy

AI's characteristics as a uniquely flexible General Purpose Technology suggest the potential for unprecedented labor market impacts, but this is not the first time society has faced significant workforce transition. Historically, labor-impacting technologies have been successfully absorbed by the economy, but the process is never painless and requires government response. The economist Carl Benedict Frey puts it well:

"Most economists will acknowledge that technological progress can cause some adjustment problems in the short run. What is rarely noted is that the short run can

[2] Specifically in Figure 4.
[3] The "reinstatement effect," coined by Daron Acemoglu and Pascual Restrepo, describes how technological automation that displaces human labor in certain tasks simultaneously creates new task categories where humans maintain a comparative advantage, thereby preserving or even increasing labor demand.
[4] To understand the range of AI's possible future capabilities and the inherent uncertainty here, consult Convergence Analysis' Threshold 2030 report. Specifically, see "Proposed Scenarios."

be a lifetime. And ultimately, the long run depends on policy choices made in the short run."[5]

Consider the case of British skilled textile workers in the early 19th century. Unhappy with displacement at the hands of machines like the power loom and the lack of any external assistance, these workers took part in the destruction of textile factories and machines as part of the Luddite movement. The rebellion was crushed, and weavers' wages sharply dropped, remaining depressed for at least the next half-century (p. 608).

In the first half of the 20th century, the U.S. demonstrated a different response to large-scale workforce changes after the Great Depression. New Deal programs like unemployment insurance and the Works Progress Administration (WPA) provided direct aid and jobs respectively to those put out of work. The WPA alone cost ~$11 billion – or ~$248 billion today – but helped save over 8 million people from the worst effects of unemployment during the Depression. Around a decade later, Congress took bold proactive action to address a massive influx of workers into the job market after World War II with the G.I. Bill. The government provided housing, unemployment insurance, and a college education to nearly eight million veterans, distributing the equivalent of around $170 billion by 1956 and nearly entirely avoiding increased unemployment after the war.

If AI truly holds the potential to automate most human tasks, not just once but continuously due to its plasticity, policies of equal or greater ambition than these historical precedents will be needed. With such a high potential for impact, it is worth overviewing what ideas policy and economics researchers have proposed to meet the moment.

---

[5] Frey, C. B. (2019). The Technology Trap: Capital, Labor, and Power in the Age of Automation. Princeton University Press. p. 18.

The Current AI Workforce Policy Landscape

Building Blocks for a Bolder Vision

AI's potential labor force implications have long been discussed in research circles, which has in turn generated a range of proposals to navigate the technology's impact. What follows is a catalog of ten interesting policy proposals currently discussed – moving broadly from more incremental to transformational. Four promising ideas are highlighted at the top, which are both highly politically feasible and useful for mitigating some of the labor impacts of AI.

The highlighted proposals are valuable, but all assume incremental and/or limited AI adoption. They are a good starting point, but if lawmakers are to plan for the case where these technologists are correct, the policy community will need to think much bigger. Moving forward, our government will likely need to take an all-of-the-above approach in order to maximize optionality and prepare for more drastic change, so more innovative policy proposals are needed.

Workforce & AI Policies: High Feasibility

Scale: Implementation Complexity & Political Barriers:
●○○○○ = low complexity/few barriers
●●●●● = high complexity/many barriers

Scale: Cost
$ = low cost
$$$$$ = high cost

Table 3

Policy Name	Description	Considerations
Create New Up-/Reskilling Programs ●●●○○ $$$	Invest in programs to educate workers displaced and at risk from the AI-economy. Includes programs to teach workers new skills and to completely retrain workers, emphasizing automation-resilient careers like many in cybersecurity and climate-tech. This support could take the form of direct grants to community colleges for displaced worker programs, tax incentives for employers to invest in training & apprenticeships for their employees, or providing expertise and	• Has instantiations that are broadly politically palatable • Hard to measure these programs' effectiveness – certain instantiations likely to have little impact • Historical analogue with the GI Bill • Challenging to predict automation-resilient skills • Less effective if many sectors are hit simultaneously or if AI's development timeline is very rapid

	capacity-building to help states design effective programs, etc.	• Success often relies on complimentary programs like income support
Modernize Unemployment Insurance ●●●○○ $$$	Modernize unemployment insurance in response to the growing threat of widespread automation. Mechanisms could include: extending benefit duration, raising wage replacement rates (~50% → ~75%), expanding eligibility to include things like partial displacement, etc. Increasing FUTA/SUTA rates – to pay for these modernizations – would also likely require complimentary taxes on labor-replacing capital to avoid further incentivizing automation. UI could also be linked to up/reskilling programs (p. 17).	• Leverages existing UI infrastructure, increasing administrative feasibility • Good opportunity for the states to test different modernization options • Potential reliance on automation taxes makes this more complicated • Also treats automation's impacts as a short-term problem; UI's effectiveness is limited in futures where this becomes a structural issue
Support Vocational Schooling ●●○○○ $$	Support educational pathways for vocational work, specifically emphasizing automation-resilient jobs (p. 21) such as plumbing and construction. Unlike upskilling, this proactively creates entirely new educational pathways for emerging workers. Support includes the government working with industry to develop curricula and credentials, modernize apprenticeships, and fund vocational academies.	• Broadly politically appealing and relatively cost-effective • More equitable and accessible than four-year degrees • Possible to misidentify automation-resilient skills • Requires a significant shift in cultural perception of vocational education • Again, only effective in scenarios where AI does not replace these types of vocational work
Establish Tax-Advantaged Lifelong Learning Accounts ●●○○○ $$	Establish tax-advantaged savings accounts for education and reskilling, similar to existing 529 education savings plans or HSA plans. Employer contributions would be tax-deductible and the government could offer matching.	• Politically palatable and implementable through existing financial infrastructure • Lower impact in transformative AI scenarios where training is not sufficient • Possibly could exacerbate inequality if only well-resourced workers can afford to contribute

Evaluation of Additional Proposals

Scale: Implementation Complexity & Political Barriers:
●○○○○ = low complexity/few barriers
●●●●● = high complexity/many barriers

Scale: Cost
$ = low cost
$$$$$ = high cost

Table 4

Policy Name	Description	Considerations
Establish a Data Dividend to Share AI Profits ●●●●○ $$ - $$$$ (for payments)	Require companies that significantly profit from AI to contribute to a sovereign fund, with proceeds supplied to citizens as a "data dividend" at a rate proportional to their data contribution. Implementation steps could include: a framework for evaluating compensation rates, reporting requirements for in-scope companies, and a public trust to manage contributions and payments. Based on longstanding ideas to compensate people for their contributions to the digital economy and existing dividends like Alaska's Permanent Fund Dividend for oil profit redistribution.	• Creates a public stake in AI without directly impacting its development • Challenging to evaluate individual data contributions and therefore implement • Would require impacting a significant portion of tech revenue to generate sizeable payments • Might incentivize companies to simply move operations – requires international cooperation to prevent this
Reduce Standard Working Hours Without a Reduction in Pay ●●●●○ $$	The current standard of 40 hours per workweek could be reduced without reducing worker pay. If AI reduces the number of total labor hours needed in the economy, reducing individual working hours to match would retain similar employment levels. Mechanisms could include providing tax incentives to adopting companies or amending the Fair Labor Standards Act to lower the overtime threshold to 32 hours a week.	• Continues the historical tradition of lowering the standard workweek as a response to increased productivity • Could further adjust as AI-driven productivity increases • Likely to see significant pushback from businesses • Realistically very difficult to manage wage impacts
Implement an Automation Tax ●●●●○ $$	Tax companies for employing labor-replacing technologies. This could take several forms, including: an equivalent tax to payroll taxes on labor-replacing capital, direct taxes on the use of AI models or robotics, taxes on increased productivity rates at high-automation firms, or high	• Directly redistributes gains from automation • Provides funding for other support programs • Likely to impact beneficial innovation

	taxes on excess profits of highly-automated companies.	• Very challenging to implement – often hard to tell exactly when a technology replaced a worker • Likely to incentivize companies to relocate – requires international cooperation to prevent this
Create Government Job Guarantee Programs ●●●●● $$$$$	The government could commit to providing employment opportunities for anyone willing and able to work, primarily focused on community needs, infrastructure, care work, and environmental restoration. Inspired by historical programs like the WPA and the PWA, the notion is to make the government an employer of last resort for displaced workers who, for one reason or another, are not part of retraining or education programs.	• Highly complicated to implement and manage, likely requiring a novel federal agency to manage partnerships with local communities • Creates societally valuable work that is not necessarily incentivized by market conditions • Extremely expensive – would likely require significant tax reform, possibly including novel automation taxes
Provide a Participation Income ●●●●● **$$$$$**	Establish a payment system providing income to citizens who participate in societally valuable activities – a "participation income" – treading a middle ground between traditional income and unconditional UBI. Activities could include: formal employment, education, caregiving, volunteering, environmental conservation, art, etc. Implementation would likely require establishing a new federal agency to define qualifying activities – in partnership with local communities – and verifying and distributing payments.	• Extremely costly and complicated to craft/implement/administer • Possible to start with pilots • Proportionally bold approach to AI futures that cause widespread long-term unemployment • Likely more politically palatable than UBI – less likely to be seen as only an entitlement • Emphasizes and supports particularly "human" tasks such as care work
Implement a Universal Basic Income ●●●●● **$$$$$**	Implement a Universal Basic Income (UBI), providing regular cash payments to all citizens without question. This would create an economic floor as AI reduces more or eventually all human labor. Such a program could be administered through a new federal agency or an expanded duty of the Social Security Administration.	• Most direct solution to structural unemployment • Unparalleled financial cost – significant concerns around inflation and necessary tax reform • Well outside of the Overton window politically • Risks a widespread crisis of purpose

The policies outlined above provide a range of tools for addressing AI's potential labor market disruption. It is likely that no single approach can fully address the unprecedented challenges posed by AI's unique high pace, pervasiveness, and plasticity, and that some combination of policy tools will be required to address workforce changes.

Lower-cost, high-feasibility approaches can serve as a first step to addressing the near-term labor market impacts of AI. However, if leading technologists' predictions prove accurate and AI capabilities rapidly advance to most economically valuable tasks, we will need simultaneous implementation of multiple approaches, including more transformational ideas. The time to develop these proposals – as well as other bold new ideas – is now, as this transformation may unfold in years rather than decades. Bridging the gap between technologists' bold predictions and economists' gradualist assumptions requires policy innovation commensurate with the potential scale of disruption.

APPENDIX

The DGPT Framework Applied to Historically Significant Technologies

Technology	Pace	Pervasiveness	Plasticity
Steam Engine	**MBWA: ~70 years** First successfully introduced in 1712 by Thomas Newcomen, but only spread rapidly in the 1800s after James Watts' improved engine, patented in 1769. Britain's stationary steam horsepower grew by a factor of 32 from 1760 to 1830, then by a factor of 60.4 from 1830 to 1907, shortly whereafter the Internal Combustion Engine (ICE) started to take over.	For decades, engineering challenges relegated the steam engine's use to pumping raw materials. Eventually it became applied to many domains, with further innovation catalyzing application in boats, rails, mills, and factories by the early 1800s.	While its number of domains handled expanded dramatically, the engine remained fundamentally a source of power. The number of tasks it handled stayed relatively constant throughout its relevance.
Spinning Jenny	**MBWA: ~20 years** One of the most impactful technologies in Britain's textile industry, the jenny had its entire rise and fall in the 40 years starting around the mid 1760s. The jenny was adopted relatively quickly, going from zero in 1763 to between 8,000-20,000 by the late 1780s. During its life, the jenny replaced numerous skilled weavers with cheaper machine operators before itself being replaced by the spinning mule and the power loom.	The jenny's main impact was in textiles, although this was a crucial industry in Britain at the time. Perhaps more importantly, the jenny was also key in launching the modern factory, starting as a (literal) cottage industry, but quickly needing dedicated spaces as the machines scaled up from 8 spindles to over 100.	The jenny was always a single purpose technology, although it remained a component of its successors like the spinning mule.
Steel	**MBWA: ~15 years** Pre-1850, steel was costly and only made in small batches (~2% of Britain's iron output at the time went to steel). That changed after Henry Bessemer's eponymous process was created in 1856. The Bessemer process spread to America in just over a decade and increased global steel production by over a factor of 50 at the turn of the century.	Steel's high tensile strength made it incredibly versatile, permeating almost every industrial domain. In the ~170 years since its first mass production, steel has seen use in weapons, trains and railroads, ships, furniture, cars, appliances, and more. It is perhaps most known for its pioneering use in massive bridges and skyscrapers.	Novel production processes like the electric arc furnace and the introduction of new alloys have allowed steel to continue to serve novel purposes nearly two centuries after the Bessemer process. Alloys like electrical steel have put steel at the center of a number of modern technologies such as motors and generators, although its use remains primarily as a foundational structural component.

Electric Motor (as a proxy for modern electric power)	**MBWA: ~30 years** The first practical DC motor was introduced by Frank J. Sprague in 1886, but around the turn of the century electricity only provided 5% of industrial power in the US. Part of the delay here is due to infrastructure, which simplified and spread in the 1910s thanks to central stations. Afterwards, electrification exploded, with power consumption in the US growing by over a factor of 20 between 1900 and 1920.	While the shift from steam power to electric power in factories was not quick, it was widespread. Moving to power-over-wires eventually revolutionized nearly every aspect of most industries and everyday life. It extended the productive hours of the day, enabled numerous new factory locations, became a source of power for modern devices, and through cascading innovation paved the way for future innovations like the internet.	Despite electricity's almost incalculable use cases, the electric motor (and the fundamental force of electricity) has stayed quite consistent. While humans have created numerous complementary inventions to harness electricity, it has fundamentally fixed principles.
The Internet	**MBWA: ~12 years** While internet experiments trace back to the 1960s, we will use ARPANET's 1983 adoption of TCP/IP as the jumping off point for viable widespread use. Only eight years later in 1991, the World Wide Web was released. In 1995, the web neared 40 million users and by 2000 it had over 400 million. One reason for this short adoption timeline was that most people already owned personal computers, but another was the High Performance Computing Act of 1991 (HPCA), which led to the creation of things like advanced fiber optic networks.	Research institutions and financial services were early adopters of the internet, even before the World Wide Web. Once the technology achieved widespread adoption, it quickly made inroads in e-commerce and online banking and saw the creation of new industries like social media. While many industries were already using the internet for things like information retrieval, the 2010s saw the digitization of many more industries like manufacturing (with IoT technology) and healthcare (with the rise of electronic health records).	The internet has changed significantly over time, but in a way that required considerable human oversight and intervention. It integrated new and evolving protocols to support things like streaming and real-time applications and its architecture now covers client-side, cloud, and edge applications. Applications on the internet also demonstrate fundamentally reorganizing principles, for example how the internet was near-universally centered around the search engine for years before (debatably) becoming centered around social media.
Generative AI	**MBWA: ~5 years** Unlike other technologies on this list, generative AI is likely early in both its development and adoption curves, but its rise has been rapid. In 2017, researchers at Google introduced the Transformer architecture, which enabled Gen AI's main consumer breakthrough, ChatGPT, in 2022. The app reached 100 million users in only two months. Since then, enterprise adoption has also exploded, with one July	Generative AI has demonstrated transformative potential in a massive range of industries from healthcare to energy to education, but what is unique is its potential to automate both cognitive and physical tasks. Historically, GPTs have primarily had labor impacts on either physical or cognitive tasks, and crossover displacement has occurred as the result of complementary innovations (consider electricity impacting physical	Unlike previous technologies, generative AI adapts to new tasks with minimal human intervention. Historical innovations required entirely new engineering for new use cases, such as repurposing steam technology from stationary engines to locomotives. Modern AI systems in contrast can be repurposed through lightweight processes like fine-tuning or prompt engineering.

	2024 survey indicating 71% of organizations use Gen AI in at least one function. Notably, this wave has not happened in a vacuum, with precursor work like the deep neural network evolving over many decades.	tasks, but being clearly foundational for the development of Gen AI). Gen AI breaks this pattern. For example, it demonstrates the potential to displace physical work when applied to robotics and to displace cognitive work when applied to digital agents.	While humans still guide this adaptation, it requires substantially less effort than retooling previous technology. The truly revolutionary aspect is AI's theoretical capacity for autonomous adaptation. As systems advance toward general intelligence, they may master novel tasks independently and even recursively self-improve. This would enable AI to automate even the new roles its introduction creates, fundamentally changing the traditional pattern of technological displacement and job creation.