Proceedings of the 5th International Conference on Interdisciplinary Humanities and Communication Studies
DOI: 10.54254/2753-7064/2026.38517

Who Shapes AI? Technology, Institutions, and Power in the Development of Artificial Intelligence

Chang Huang

International Department, Affiliated High School of South China Normal University, Guangzhou, China
jasonhuang1234@qq.com

Abstract. Artificial intelligence is usually mainly explained through technological breakthroughs, such as larger models, improved algorithms, stronger computing power and expanded datasets. This article challenges this interpretation by viewing artificial intelligence as a sociotechnical system whose development and consequences are constantly influenced by social institutions. Based on research of the history of the Internet and human-computer interaction, institutional development, automated responsibility, and algorithmic labor, this paper compares early computing systems with contemporary artificial intelligence. Analysis shows that technical capabilities themselves cannot determine how a technology develops or how it affects society. Governments, universities, companies, markets, legal systems, workers and users all influence the way artificial intelligence is funded, deployed, regulated and understood. Meanwhile, these actors do not have equal power. Cases such as moral collapse zones and algorithmic wage discrimination demonstrate how automation unevenly redistributes responsibility, information and economic benefits. Therefore, this article holds that the future of artificial intelligence is not determined by the technology itself, but rather by the institutional structure and power relations in which the technology is developed and utilized.

Keywords: Artificial intelligence, larger models, computing systems

1. Introduction

Artificial Intelligence is often described as a technological breakthrough. Larger model, faster processors, better algorithms and greater datasets seem to explain the reason why AI has developed so rapidly. This explanation appears convincing since it provides a clear source for the technological change: As scientists and engineers create increasingly capable machines and algorithms, society adapts to what those AI can do and develops. However, this explanation mistakes one part of technological development for the whole process. A technically capable system does not automatically become a socially significant technology. Research and experiment need funding support, infrastructure needs to be constructed, technical standards need to be chosen, products need markets, users need to adopt the new technology, and governments also need to decide which kind of usage are permitted. The algorithms and capabilities of the new technology is essential, but so do the institutions that determine how those capabilities are developed and applied.

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2753-7064/2026.38517

The history of computer development makes this distinction particularly important. For instance, the Internet is not an inevitable outcome of any technological invention. It has developed in the process of continuous interaction among government agencies, universities, engineers, standard organizations, private enterprises and users. Janet Abbate's research on the history of the Internet particularly emphasizes these social and cultural conditions, pointing out that the design of technology itself is formed in the cooperation and conflict among multiple institutions and groups [1]. The development of artificial intelligence is also going through a complex process. Its development direction is not only influenced by the possibilities of engineering and technology, but also by factors such as government research funding, market rewards for specific applications, enterprises' deployment of AI, the way workers are managed by AI, the allocation of responsibilities by legal systems, and society's cultural expectations for automation.

Therefore, rather than understanding artificial intelligence as an autonomous product driven solely by technological innovation, it is better to view it as a socio-technological system. This view does not deny the importance of algorithms or computing power. On the contrary, it emphasizes that technical capabilities will only acquire specific significance and consequences after entering the social system. Just like the Internet, the development of AI is continuously shaped by governments, markets, scientific research communities, users, legal systems, and cultural values. More importantly, these forces do not possess the same degree of power. When the institutions that control AI have more information and decision-making power than the people affected by AI, social shaping may produce exploitative results, not just technological progress. The concept of "moral crumple zone" proposed by Madeleine Clare Elish and Veena Dubal's analysis of "algorithmic wage discrimination" [2, 3]. This reveals such a problem: automation may redistribute responsibility and economic value, protecting powerful institutions while individual workers bear more uncertainties. Therefore, the truly important question is no longer merely what AI can do technically, but who can shape this system, based on what kind of incentive mechanism, and what the ultimate consequences will be.

2. The internet as a socio-technical model

The development history of the Internet provides a useful historical basis for this argument, as it challenges the common imagination that "technological innovation mainly comes from isolated inventors". As one of the most important predecessors of the modern Internet, ARPANET originated in a research environment supported by the US government and closely linked to universities and research institutions. Katie Hafner and Matthew Lyon described early networks as the result of collaboration among researchers, engineers, government agencies and scientific research organizations, rather than the personal achievement of an individual inventor [4]. This technology, of course, cannot do without professional technical knowledge, but technical knowledge alone is not sufficient to create a network. The structure of funds, institutional relationships, research priorities, and whether organizations are willing to cooperate are equally crucial.

Abbate further pointed out that even though the early technological foundation of the Internet has been established, it continues to change. Different groups engaged in debates around technical standards. Users have developed new uses that the designers might not have anticipated initially. Commercialization has further transformed the scale and functions of the Internet. As users themselves can also participate in the formation of new practices and applications, the boundary between producers and users has become particularly blurred. Therefore, the Internet was not "invented first and then introduced into society". On the contrary, its architecture and application

DOI: 10.54254/2753-7064/2026.38517

have gradually taken shape in the continuous interaction between technological possibilities and social choices [1].

Benjamin Peters' research on the failure of the construction of the national computer network in the Soviet Union provides a particularly valuable comparison. Soviet scientists had considerable technical capabilities and also proposed ambitious national network plans, but institutional competition and bureaucratic conflicts repeatedly hindered the true implementation of these plans. Peters contrasted this failure with the stronger research collaboration environment and national support during the development of ARPANET in the United States. This comparison reveals a broader conclusion: Possessing technical knowledge is not sufficient to guarantee the true development of technology. The institutional organization itself can determine whether a technological concept can eventually become a system in reality [5].

This historical perspective has changed the way we understand artificial intelligence. Contemporary narratives about AI often organize its development history as "technological breakthroughs", such as advancements in neural networks, improvements in machine learning methods, more training data, and enhancements in computing power. All of these are important, but if they are regarded as complete explanations, it would repeat the problems that Abbate and other research institutes on the history of technology have attempted to correct. An AI model will only truly have social influence when relevant institutions provide development resources and organizations find ways to deploy it. The government influences research through policies and research funds. Universities and laboratories determine the priorities of scientific research. Private companies make investments based on business expectations. Users provide data, use the system for new purposes, and sometimes also object to the way the technology is deployed; The law determines which data collection, automated decision-making and labor practices are permitted.

Of course, the similarities between the Internet and artificial intelligence should not be overly exaggerated. AI is not simply a repetition of the development history of the Internet, and the technical structures of the two systems are also different. The truly valuable aspect of the history of the Internet lies in the fact that it provides an analytical model. It shows that even if a technology is highly complex in terms of technology, it still depends on the surrounding social arrangements. In this way, the focus of researching the development of AI should not merely be on asking "Who invented this system?", but rather on further inquiring: What systems make this system possible? What interest forces have influenced its development direction?

The commercial transformation of digital technology has made this institutional issue even more important. Richard Barbrook and Andy Cameron's criticism of "Californian Ideology" challenges a technological culture that combines digital liberalism, market beliefs and individual entrepreneurship. Their viewpoint indicates that the narrative about the "inherent liberating nature" of technology may mask the economic structure of its actual development. A technological system may seem decentralized and empower individuals on the surface, but infrastructure, capital and decision-making power can still be highly centralized in reality. This contradiction is even more pronounced in the AI era, as developing large-scale AI systems often requires huge amounts of resources, which are mainly in the hands of governments, top universities, and large technology companies [6].

3. Human–machine collaboration and the future of work

In his 1960 article "Man-Computer Symbiosis", J. C. R. Licklider proposed a different vision for the development of computers. Licklider did not simply describe the future development of computers as "machines gradually replacing humans", but envisioned a cooperative relationship in which

DOI: 10.54254/2753-7064/2026.38517

humans and computers contribute different capabilities respectively. Humans are still responsible for goals, assumptions, evaluation criteria and final judgments, while computers are responsible for those operations that can be programmed and automatically executed, thereby supporting human thinking and decision-making. His model is important because it shows that the relationship between humans and intelligent machines never has to be understood as pure competition [7].

This is particularly important for today's discussion on AI and employment. Public discussions often condense the debate on future jobs into one question: Will AI replace human workers? This way of asking questions tacitly assumes a premise that technical ability can directly determine the outcome of labor. If a machine can complete a certain task, then the person who originally accomplished this task will disappear. But if viewed from the perspective of the socio-technological system, the problem will become different. The significance of AI lies not only in which tasks can be automated, but also in how organizations rearrange their work around this automation.

Enterprises may use AI to reduce certain types of positions, but they may also use the same technology to enhance the productivity of existing employees, change their responsibilities, or create new forms of "human-machine collaboration". None of these results are automatically determined by the algorithm itself. They depend on labor costs, management strategies, legal regulations, employees' bargaining power, and how the productivity benefits generated by technology are ultimately distributed. Therefore, Licklider's concept of human-machine symbiosis is valuable not because it accurately predicts all the features of today's AI - when he put forward this view, it was still very far from modern machine learning and large language models - but because it indicates that "substitution" is merely a possible institutional arrangement.

The distinction between "substitution" and "reorganization" further explains why AI should be understood as a power relationship system. If employees can use AI to complete repetitive tasks more quickly while still retaining their judgment and decision-making power, then automation is closer to the cooperative relationship envisioned by Licklider. However, if employers use AI to continuously monitor their employees, determine wages through opaque formulas, or weaken employees' control over the work process, then the same broad category of technology will generate completely different social relationships. What has truly changed is not only technical capabilities, but also the redistribution of power among different individuals and institutions.

Therefore, the future of work can neither be understood completely optimistically nor by technological determinism. AI can indeed enhance human capabilities, but these benefits will not be distributed automatically and fairly. The truly crucial question is: Who controls this technology and who can reap the productivity benefits it generates. The socio-technological perspective thus transcends the simple opposition between "AI will replace everyone" and "AI will only help humans". Both of these outcomes can occur in different institutional environments because technology is always interacting with the existing economic and institutional structures.

4. Power, responsibility, and algorithmic labor

The concept of "moral crumple zone" proposed by Elish very clearly explains how institutional structure affects the consequences of automation technology. In complex automation systems, responsibilities are often scattered among designers, enterprises, technical systems, managers, and human operators. However, once a failure occurs, public and legal attention often focuses on the person closest to the system at the time of the accident. The "moral buffer zone" Elish referred to is precisely such a situation: a human actor who actually has limited control over the system bears a disproportionate moral or legal responsibility for the behavior of a highly automated system [2].

DOI: 10.54254/2753-7064/2026.38517

This concept reveals a contradiction in automated narration. When an automated system operates well, people may describe it as proof that "technology can operate independently of humans". However, when the same system malfunctions, the responsibility is suddenly shifted back onto the human operator. In other words, when "autonomy" can bring value, machines are regarded as autonomous; But when it comes to finding someone to take responsibility, human responsibility resurfaces. This contradiction cannot be simply explained through technical design, but rather stems from the institutional arrangements of the organization regarding power, responsibility and accountability mechanisms.

Therefore, the "moral buffer zone" reveals how AI as a socio-technological system could potentially become exploitative. Exploitation is not because the algorithm itself has the motive to harm people, but because the system may separate control from responsibility. Employees may be responsible for the output of a system, but in reality, they have almost no real control over the design, training data or operating environment of this system. Organizations gain the benefits brought by automation, while individuals retain or even bear more risks of failure. On the surface, this is merely a technical task allocation, but in reality, it is also a responsibility allocation in a political sense.

Dubal identified similar power imbalances in the economic organizations of platform labor. The algorithmic wage discrimination she proposed describes a system that uses highly refined data and constantly changing algorithmic formulas to set different and personalized wages for employees engaged in roughly similar labor. Compared with transparent fixed hourly wages, algorithm-determined income may make it difficult for workers to predict their own earnings and understand exactly how the platform determines remuneration [3].

The significance of these systems does not merely lie in the fact that they employ complex software. The real key lies in the asymmetry between information and institutions. The platform may have access to a vast amount of information about market demand, worker behavior, location, order-taking patterns, and other variables, while individual workers have very little knowledge of the systems that determine job opportunities or wages. Therefore, the algorithm actually operates in an unequal relationship: one party can observe, calculate and adjust the entire system, while the other party must make economic choices under the circumstances of a severe lack of information.

This also indicates that merely describing AI as a "technology for enhancing efficiency" is not sufficient. Algorithmic systems may indeed coordinate a large number of workers more efficiently, but "efficiency" cannot answer the question of "fairness". A system can collect information, assign tasks or set personalized prices very efficiently, while weakening employees' ability to understand or challenge working conditions. Therefore, Dubal's analysis links technical design to labor laws, wage regulations and bargaining power, rather than viewing algorithmic management as a neutral technological improvement.

Elish and Dubal study problems in different fields, but they both reveal similar structures. In the "moral buffer zone", institutions have more control over automated systems, while individuals with less power bear the responsibility. In algorithmic wage discrimination, the platform has more information and control over wage setting, while employees face greater income uncertainty. In both cases, automation does not eliminate social relationships but rather reorganizes them [2, 3].

This model can more specifically explain why AI has become exploitative. The problem does not lie in the fact that socio-technological systems are inherently exploitative. Licklider's model has demonstrated that computing technology can also be organized around complementarity and cooperation. What truly makes exploitation more likely to occur is when the social structure around AI concentrates control, information and economic benefits in the hands of a single institution,

DOI: 10.54254/2753-7064/2026.38517

while transferring uncertainty, risks and responsibilities to individuals with less power. The decisive factor is thus not how much "intelligence" the machine itself possesses, but rather to what extent there is power inequality in the social system around the machine.

This also helps to clarify the relationship between "technical bias" and "institutional bias". AI bias is often discussed as a problem mainly arising from data or code flaws, as if it could be resolved simply by engineers making better technical corrections. Of course, technological improvement is important, but many harmful consequences are actually related to what purpose the system is used for and in which institutional environment it operates. Even if a wage algorithm is technically perfectly normal, if its goal is to maximize the value that the platform can extract by leveraging personalized worker information, it may still have unfair consequences. Similarly, even if an automated system is technically highly accurate, if organizational rules require operators with almost no control to take responsibility, the distribution of responsibility may still be unfair. Therefore, technical performance and social equity are two distinct issues.

5. Discussion: social shaping and institutional choice

Understanding AI as a system shaped by society can also prevent technological determinism from becoming an excuse to excuse institutional choices. Statements such as "AI will inevitably eliminate these jobs" make specific economic outcomes seem inevitable. However, the development history of computing technology shows that technology rarely automatically determines a single form of social organization. It is the specific systems that truly determine where automation is used, which human jobs continue to exist, how productivity gains are distributed, and which risks are acceptable.

This does not mean that society can have unlimited control over technology. Technical capabilities themselves do indeed create real possibilities and limitations, and the system cannot simply determine that any desired technical outcome will definitely be achieved. At the same time, the socio-technological perspective does not imply that all participants have the same influence. On the contrary, one of the most significant characteristics of contemporary AI is that the distribution of resources required to shape it is highly unequal. Large enterprises may control computing infrastructure, proprietary models, user data and huge amounts of capital, which individual users or workers simply cannot compare with. Although the government can regulate these companies, it may also rely on them in terms of technical expertise or infrastructure. Researchers can influence research directions, but their work is often restricted by the priority of external funds.

Therefore, the future of AI is indeed shaped by society, but this does not mean that it is naturally "democratically" shaped by society. Recognizing the significance of society does not mean that "the entire society" will equally determine what AI will become. Different groups have different degrees of power. A more meaningful question should be: Exactly which institutions are shaping AI?

This perspective also indicates that regulation should not be added only after the technology is developed. Law itself has long been a part of the environment in which technological systems are formed. Legal rules in areas such as privacy, worker identity verification, liability, market competition, intellectual property rights, and discrimination will directly affect the incentive mechanisms that companies face when designing and deploying AI. The market does not passively wait for the technology to mature before responding. The expectation of profits itself will affect which technologies can receive investment. Users are not completely passive recipients either. Their actions may create new applications, expose system failures, or force institutions to change their technologies. Therefore, it is difficult to truly separate the technical and social aspects of AI, as the two are constantly changing each other.

DOI: 10.54254/2753-7064/2026.38517

The history of the Internet still holds value here because it simultaneously reveals possibilities and limitations. Abbate's research demonstrates how flexible network architectures, institutional collaborations, and user engagement can drive unexpected innovations; Peters, however, demonstrated the opposite possibility: even with a powerful technological concept, institutional conflicts may hinder the true realization of the technology. These histories cannot predict how AI will eventually develop, but they make one conclusion extremely hard to deny: the development of technology cannot be explained merely by technical capabilities [1, 5].

AI has further strengthened this point as it is increasingly involved in decision-making regarding work, responsibility, information and economic opportunities. The impact of AI depends on the social structure through which these decisions are made. Licklider's vision represents a possible arrangement in which humans and machines play complementary roles. Elish and Dubal, on the other hand, revealed another arrangement in which automation might deepen the existing power asymmetry. The difference between these two types of results is not merely how complex the technology is, but rather how different the systems that organize this technology are.

6. Conclusion

Artificial intelligence is often described as a force that enters society from the outside and then changes it, but in fact, the causal relationship is bidirectional. AI does indeed change systems and human behavior, but the choices made by systems and people also simultaneously determine which AI will be developed, how it will be used, and ultimately whose interests it will serve. Technology itself is real and its capabilities are indeed important, but these capabilities do not determine their social significance on their own.

The history of the Internet shows that technological systems are formed in the continuous cooperation between technological actors and social actors, rather than being created by a single isolated invention. Licklider's early vision of human-machine symbiosis further illustrates that the improvement of machine capabilities does not necessarily mean that humans will definitely be replaced. However, contemporary AI also brings new risks, as the institutions that shape it often hold unequal information and power. The moral buffer zone proposed by Elish indicates that responsibility may be transferred to those who actually have limited control. Dubal's algorithmic wage discrimination indicates that information and computing power may be used to personalize wage Settings, thereby weakening workers' certainty about their own income. These cases are not accidental consequences outside of technological development, but rather concrete outcomes formed when technology operates within specific legal, economic and organizational structures [7, 2, 3].

Therefore, the most crucial question regarding the future of AI is not whether the technology will be powerful enough to determine the future of society, but how society will organize the power that already exists around this technology. AI may automate certain tasks or eliminate certain types of work, but it can also create new tasks and deepen the cooperation between humans and computers. Which outcome ultimately dominates depends on a series of social choices such as investment, regulation, labor organization, institutional accountability, and how the benefits of technology are distributed.

Therefore, artificial intelligence should neither be regarded as an unstoppable and completely independent force, nor as a completely neutral tool whose consequences depend only on how an individual uses it. AI is a socio-technological system that has emerged in the continuous interaction between technological capabilities and unequal social systems. Recognizing this relationship will make the future of AI seem less "predetermined", but at the same time, it also means that society must take more responsibility for the final outcome.

DOI: 10.54254/2753-7064/2026.38517

References

[1] Abbate, J. (1999). *Inventing the Internet*. MIT Press. Cambridge, MA.
[2] Elish, M. C. (2019). Moral crumple zones: Cautionary tales in human-robot interaction. *Engaging Science, Technology, and Society, 5*, 40–60.
[3] Dubal, V. (2023). On algorithmic wage discrimination. *Columbia Law Review, 123*(7), 1929–1992.
[4] Hafner, K., & Lyon, M. (1996). *Where wizards stay up late: The origins of the Internet*. Simon & Schuster. New York.
[5] Peters, B. (2016). *How not to network a nation: The uneasy history of the Soviet internet*. MIT Press. Cambridge, MA.
[6] Barbrook, R., & Cameron, A. (1996). The Californian ideology. *Science as Culture, 6*(1), 44–72.
[7] Licklider, J. C. R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics, HFE-1*(1), 4–11.