Ethics of AI in Lift Systems

Matthew Appleby

Peters Research Ltd, UK

Keywords: artificial intelligence, ethics, lift systems, fairness, privacy, accountability, ownership

Abstract. Artificial intelligence is increasingly being integrated into lift systems, with applications in dispatching, traffic analysis, preventive maintenance and passenger-facing services. These technologies may improve efficiency, reliability, safety, accessibility and sustainability, but they also introduce ethical risks that have not yet been widely explored within lift research.

This paper argues that lift systems represent a distinct ethical context. They are shared, physical, safety-critical infrastructure embedded within buildings, and many users can't meaningfully opt out of using them. The ethical burden is therefore different from many consumer digital systems, where individual choice and consent are often treated as central safeguards.

As lift systems become more automated and data-driven, questions arise about prioritisation, surveillance, labour, environmental cost, commercial incentives, accountability and control. The paper proposes early guidelines intended to support industry discussion, future standards development and ethically aware technical research. The aim is to ensure that deployment of AI improves the systems people depend on while promoting fairness, equality, privacy, skills use, sustainability and accountability.

1 INTRODUCTION

1.1 Background

In 2025, the author published a technical paper which included a literature review of previously published research on the subject of AI in lift systems [1]. The paper classifies existing research by use case and speculates on future developments that could emerge as technology advances.

The paper includes a brief section on the ethical concerns accompanying these technological developments, concluding that "the industry has an ethical responsibility to care for its workers and passengers as well as the environment". It is noticeable how few papers cited in the literature review explored the ethical ramifications of the technological progress described in the research.

Unlike the previous research, this is not a technical paper. This is an ethics paper which explores the issues the lift industry will face as AI and automation continue to be introduced into systems that move people through buildings.

1.2 Scope

This paper focuses on AI, but some of the issues discussed also arise from automation more generally. A conventional rule-based controller can still raise ethical questions if it prioritises some passengers over others. A monitoring system can create privacy concerns even if it does not use machine learning. This wider scope is useful because AI is difficult to define precisely, and because many real systems combine statistical models, rules, sensors, cloud services and human procedures.

The paper does not provide a technical background on AI or the lift industry, as this was covered in the previous paper [1].

1.3 Ethical Framing

In many digital contexts, consent and choice are treated as central safeguards. Lift systems are shared, physical, safety-critical infrastructure embedded in buildings that many users can't meaningfully opt out of. A person with impaired mobility or a resident in a high-rise building may be unable to choose the stairs. This lack of meaningful exit changes the ethical burden. The focus must therefore shift from individual choice toward system-level responsibility.

2 ETHICAL RISKS

This section identifies seven ethical risk topics. The categories are not exhaustive, but they provide a practical structure for discussing how AI-enabled lift systems may cause harm if deployed without suitable safeguards.

2.1 Fairness

Lift service is shared, which means each dispatching decision can impact everyone in the system. A dispatching decision that improves one passenger's journey can worsen another's. For example, reducing the number of stops on someone's journey might shorten their transit time but increase another's waiting time. Prioritisation is therefore unavoidable in any dispatcher-based lift system. The ethical issue isn't whether lift systems prioritise, but which forms of prioritisation are justified and who decides. Priority for disabled users may improve fairness by compensating for a disadvantage. For example, a destination control system that monitors passengers and allows extra dwell time for a visually impaired passenger might improve accessibility without separating that passenger from other passengers, thereby making the lift system more inclusive. Paid priority, on the other hand, may reduce fairness by giving better service to those with greater economic power. A system that consistently provides premium tenants with better service may degrade the ordinary service available to everyone else and reduce the overall performance of the lift system.

In lift systems, a technically efficient outcome may still be unfair if it systematically disadvantages people who are already less able to use the building. Fairness can't be reduced to a single numerical target such as average waiting time but should instead consider context and the variance of passengers [2].

2.2 Privacy

AI-enabled lift systems may require data about passengers, lift movement and building use. Some data is low risk, such as aggregate counts of calls or stops. Other data can be highly sensitive, particularly where cameras, biometric systems or behavioural profiles are involved. Research on smart buildings shows that even apparently innocuous data, such as information about occupancy, can reveal private patterns when combined with other data [3].

A camera installed for passenger counting may also be capable of identifying employees, tracking visitors, analysing behaviour or supporting targeted advertising. AI ethics work on smart homes similarly identifies privacy, autonomy, transparency and fairness as key issues where AI systems observe domestic or semi-domestic life [4].

Advertising in lifts is not merely a question of visual annoyance. When combined with analytics, it can become a reason to identify, classify and monetise passengers in a space they can't meaningfully avoid. A phone user may be able to delete an application or change a setting, but a lift passenger can't change the data policy of the lift car.

2.3 Labour

AI-enabled preventive maintenance may reduce unnecessary inspections and help technicians arrive with the right parts, tools and information. This is a genuine benefit as it means lifts spend less time out of service, breakdowns may be prevented, repairs will be less error-prone, and technicians can focus on work that is needed rather than troubleshooting.

However, data-driven maintenance also changes the nature of work. If technicians are given a list of AI-generated tasks, their role may shift from diagnosis to execution. The skill required to interpret symptoms, understand unusual failures, and reason about the whole system may decline [5]. This risk is related to a well-known irony of automation: automated systems may remove routine human practice while still expecting humans to intervene successfully when something unusual happens [6].

Work changing because of AI isn't a problem in isolation; all technology changes work. The problems occur when workers aren't given the power, training and security to shape that change and when efficiency gains are captured elsewhere while workers bear the disruption. If expertise is allowed to decay, the industry may also become less resilient when automated diagnostics are wrong, incomplete or unavailable.

2.4 Sustainability

AI could improve sustainability in lift systems: better dispatching may reduce unnecessary travel, predictive maintenance may reduce avoidable component replacement and emergency journeys, and better expert design may reduce over-provision in buildings [1]. These are not trivial benefits, particularly in a sector linked to buildings and infrastructure [7].

At the same time, AI has environmental costs. AI depends on data centres, electricity, cooling, hardware manufacture and network infrastructure. The European Parliamentary Research Service identifies energy demand and environmental impact as gaps in many AI ethics frameworks [8]. The International Energy Agency reports that data centre electricity demand is projected to grow substantially as AI expands [9]. Recent work on water use also notes that AI and data centres can place pressure on water systems through cooling and electricity generation [10].

There is a risk that the lift industry will implement new systems that it sells as sustainable when it's actually just outsourcing the environmental impact to a different industry sector.

2.5 Incentives

Many ethical failures do not begin with malicious intent. They emerge when incentives reward behaviour that slowly degrades the system for users. Cory Doctorow uses the term 'enshittification' to describe the process by which digital platforms initially serve users well, then increasingly optimise for business customers and finally for the platform itself [11]. The word is deliberately abrasive, but useful because it captures a pattern of service degradation driven by commercial incentives, not bad actors.

A comparable risk could emerge in AI-enabled lift systems. A building may first use data collection and AI to improve passenger service; for example, a dispatcher may use facial recognition technology to anticipate a passenger's calls before they put them into the system. Later, the data from that system could be sold to a micro-managing employer tracking their employees' every movement or to an advertising agency that can target passengers against their will.

It's important to emphasise that this pattern should not be seen as a moral failing of the companies involved, but as a detrimental pathway that industries take when the economic system incentivises it.

2.6 Accountability

Accountability must be distinguished from blame. Blame asks who should be punished after a failure, while accountability asks who can explain the failure and how the system should change to prevent it from recurring.

Decisions about lift systems are distributed between many parties, including controller manufacturers, software suppliers, building owners, maintenance providers, consultants and operators. AI-enabled systems may spread this distribution further, making it harder to identify who is answerable when something goes wrong and increasing the risk that responsibility becomes obscured.

This risk is reinforced by organisational fragmentation. A lift system may involve multiple organisations, each with only partial control, making it possible for responsibility to be passed around the network. In some cases, harm may emerge not from a single negligent decision, but from the interaction of many reasonable local decisions. This can create accountability gaps which often lead to responsibility being misattributed.

One tendency is to shift responsibility onto the algorithm itself. If a dispatching system, maintenance model or passenger-monitoring system produces a harmful outcome, it may be said that "the AI decided". This framing is misleading as an algorithm is not an independent moral actor. Treating it as the responsible party risks turning technical opacity into institutional evasion.

The opposite tendency is to blame the nearest human. Madeleine Clare Elish's concept of the moral crumple zone is useful here. A moral crumple zone occurs when a human actor absorbs blame for the failure of a complex automated system, despite having limited control over it [12]. In lift systems, a technician, building manager or call-centre operator may be close enough to a failure to be blamed, but not powerful enough to have shaped the design, training data, commercial objectives or operating rules that produced it.

Whilst organisations who use fragmentation to obscure responsibility should be held accountable, ethical accountability should not require inventing fault where none exists. Some failures may be genuinely unforeseeable.

2.7 Ownership

The OECD Recommendation on Artificial Intelligence lists "transparency and explainability" as a core principle [13]. This paper argues that transparency is not sufficient and instead includes it as a prerequisite for ownership.

If passengers, workers and building users are affected by AI decisions but have no influence over the system's goals, transparency alone can't solve the ethical problem. Many of the risks identified in this section depend on who controls system goals, data and decision-making.

Ownership also concerns the means of computation. Lift companies are unlikely to start running self-hosted cloud platforms, meaning the lift industry will ultimately become dependent on a handful of cloud platform providers. Ethical decisions about access, prioritisation and surveillance may be made far from the people impacted by them. A future ethical framework for lift AI must therefore consider not only what systems disclose, but who has the power to shape them.

3 ETHICAL GUIDELINES

The following guidelines are offered as a starting point for industry discussion and future standards development. Their purpose is to respond to the risks identified above by suggesting practical directions.

3.1 Fairness

Prioritisation can't be eliminated from lift dispatching, nor should that be the goal of a fair dispatcher. However, the basis of prioritisation should be made clear during design and procurement rather than hidden inside an optimisation function, and the target should always be to use prioritisation to rebalance existing inequalities.

A useful starting point is to distinguish between prioritisation that compensates for disadvantage and prioritisation that reinforces advantage. Accessibility-related priority may be justified when it helps users access a building on more equal terms, whereas paid priority warrants stronger scrutiny because it may improve service for one group by degrading service for others.

Fairness could be expressed through constraints as well as optimisation targets. Minimum service levels, maximum disadvantage limits and accessibility constraints may be needed to prevent efficient systems from producing unfair patterns. In practice, this may mean that a dispatcher is allowed to optimise performance only within certain boundaries.

3.2 Privacy

AI-enabled lift systems should collect the least identifiable data necessary for the task. The ethical default should be to measure flows rather than identify people. Passenger counts, call data and car movements are often sufficient to improve dispatching or maintenance without collecting images, biometric data or individual behavioural histories.

There are two layers of limiting data collection. The first is to use low-resolution sensors such as light curtains, accelerometers and load sensors rather than cameras, biometrics and mobile-device tracking. This layer reduces the data points that can be misused. The second layer is to process the raw data close to the source and only store aggregated data in the cloud. Where cameras are necessary for counting passengers, systems should process video frames locally and upload passenger counts to the cloud, rather than storing identifiable footage in the cloud for processing there. In these cases where more intrusive sensing is proposed, the purpose should be narrow and justified.

3.3 Labour

AI should support expertise rather than quietly remove it. Predictive maintenance tools can improve lift reliability and help technicians work more effectively, but they should not reduce skilled workers to passive executors of instructions. A responsible system should explain why a task has been recommended, what evidence supports it and when the recommendation may be unreliable. Used well, AI could become a tool for continuing professional development, helping technicians understand patterns across many installations rather than simply telling them what to replace.

If technicians are expected to act on automated diagnostics, they should also be able to challenge, correct and improve those diagnostics. Their feedback should be treated as valuable training data, not as resistance to automation. This would make AI systems safer while preserving practical expertise and resilience when automated diagnostics are wrong, incomplete or unavailable.

However, participation and training do not fully answer the displacement problem. If AI-enabled maintenance reduces the labour required to maintain a portfolio of lifts, some workers may lose work. If AI-integrated robotics eventually becomes capable of performing more physical maintenance tasks, even more workers could be at risk of unemployment.

A responsible transition, therefore, requires serious industry discussion about retraining, redeployment, workload reduction and how efficiency gains should be shared. Otherwise, workers may be asked to help build systems that reduce their own bargaining power or employment security.

3.4 Sustainability

AI should be deployed where it improves total system sustainability, not merely local performance. A system that reduces lift energy consumption but requires disproportionate computation, hardware, cloud storage or sensor replacements may not be environmentally beneficial overall.

This requires comparison between possible implementations. AI should not be used simply because it is fashionable or technically possible. Designers should compare AI models, simpler automation, rule-based logic, calculation, simulation and conventional engineering improvements, considering both performance and environmental cost. Resource-intensive AI is more ethically defensible where it enables benefits that simpler methods can't provide, such as early fault detection from complex sensor data or adaptive accessibility functions. For example, if early fault detection reduces component replacement, emergency callouts and premature modernisation, it may be environmentally preferable even after accounting for the sensors, processing and data infrastructure required.

The aim should be to make the environmental cost of digital infrastructure visible in design decisions. If two approaches deliver similar benefits, the less resource-intensive approach should normally be preferred.

3.5 Incentives

If AI-enabled lift systems create new revenue streams through advertising, paid priority, or behavioural analytics, there will inevitably be pressure to optimise around those revenue streams at the expense of service quality. An ethical approach to incentives should therefore preserve competition, substitution and interoperability.

A building owner should not become dependent on a supplier because historic maintenance data, traffic data, diagnostic models or optimisation settings can't be exported or interpreted elsewhere. Open-source interfaces that allow interoperability and competition reduce the likelihood of enshittification. If a system is genuinely good, it should not need to rely on making exit impossible.

Some of these issues may ultimately require competition law or changes to existing rules that prevent interoperability, but the lift industry does not need to wait for a complete legal solution. Industry bodies and standards groups could require AI-enabled lift systems to meet basic anti-lock-in conditions. These might include documented interfaces, data portability, the ability to change analytics providers without losing essential records, and a clear separation between safety-critical functions and optional commercial services.

A successful ethical framework should not treat profit as the problem. The problem is a market structure in which profit can be increased by degrading service, restricting interoperability or trapping customers. The aim should be to create a system where companies make more money by maintaining high-quality, interoperable and long-lived systems than they would by lowering quality after customers are locked in.

3.6 Accountability

Accountability should make responsibility traceable without reducing every failure to blame. A fair accountability system must avoid both responsibility evasion and scapegoating. This does not mean that every failure must have a single guilty party, but it does mean that failures should not be dissolved into technical or organisational complexity.

The aim is to create systems that are auditable, explainable and learnable. When something goes wrong, the industry should be able to identify whether the problem was caused by poor design, poor configuration, inadequate maintenance, unforeseeable circumstances or some interaction between

them. Accountability is not only about assigning responsibility after failure, but about ensuring that responsibility is visible before failure occurs.

3.7 Ownership

Ownership, in this context, should be understood as control over the system's purpose and direction, not merely possession of hardware or data. As AI becomes more embedded in lift systems, the industry's domain expertise should not be displaced by external cloud service providers whose priorities may be driven more by platform growth than by lift performance.

Democratisation does not mean every passenger configures the dispatcher, but that the values embedded in the system should be shaped by those with a legitimate stake in its operation. People such as building owners, maintainers, consultants, manufacturers, passengers and workers may not all need the same level of control, but they should not become passive subjects of decisions made entirely elsewhere. The choices that shape passenger experience and building operation should remain close to the lift system and the people affected by it.

A future ethical framework should ensure that the lift industry and its users retain meaningful power over what AI-enabled lift systems are for.

4 DISCUSSION

4.1 Possible futures

The same technologies could produce very different futures. In an optimistic future, AI-enabled lift systems improve accessibility, reduce unnecessary waiting, prevent breakdowns, support technicians and reduce waste. A visually impaired passenger might be guided to the right car without needing to identify themselves publicly. A technician might be sent to site with the correct diagnosis, tools and parts before a failure occurs. A building designer might use better traffic data to avoid both under-provision and unnecessary over-provision.

In a pessimistic future, the same technologies could enable surveillance, targeted advertising, paid priority, workforce deskilling and responsibility avoidance. A system originally justified by safety or service improvement might later be used for behavioural profiling. A dispatcher designed to improve flow might be tuned towards commercial priority. A maintenance system designed to support technicians might become a means of reducing skill, headcount and accountability.

The difference between these futures will not be determined by the technology. It will depend on governance, ownership, incentives and the industry's willingness to treat ethical design as part of engineering design.

4.2 Further research

This paper has attempted to organise the ethical problem space rather than solve it. Further research is needed to translate these early guidelines into practical tools for the lift industry.

The first area is measurement. The industry needs better ways to evaluate AI-enabled lift systems beyond average waiting time, local energy use or maintenance costs. Future work should investigate how fairness, privacy risk, environmental impact, resilience, accessibility and long-term service quality can be measured in ways that are practical for design, procurement and auditing. These measures will not remove the need for human judgement, but they may make ethical trade-offs easier to see and discuss.

A second area is governance. Further research should examine which issues are best addressed by changes to the law, technical standards, procurement requirements, industry body approval or professional guidance. Some problems, such as anti-competitive lock-in or misuse of personal data,

may require legal intervention. Others may be better addressed through industry rules and standards of good practice. The lift industry should not wait for a perfect regulatory framework before beginning this work.

A third area is participation. If AI-enabled lift systems impact passengers, workers, building owners and maintainers, then the industry needs practical ways for those groups to influence how systems are designed and operated. This includes research into how expertise from technicians and consultants can be preserved, how passengers' interests can be represented, and how control over data and computation can remain close to the lift system rather than drifting entirely into external platforms.

Finally, ethics should not be confined to papers specifically about ethics. Often, those closest to a technical project are best placed to understand its ethical implications. Future lift research that develops the technological body of knowledge should therefore consider the ethical consequences of the work.

4.3 Declaration of AI use

In light of the transparency and accountability I have been espousing throughout this paper, it would feel wrong not to disclose my use of AI in conducting this research. Many of the sources cited in this paper were found by a large language model (LLM). I reviewed all the sources before uploading them to a new LLM instance, which was used to compose draft sections of the paper.

My previous paper on AI in lift systems was written in much the same manner.

The use of AI to assist with this research does not transfer any responsibility away from me. I take full responsibility for this paper, including the themes, narratives, conclusions and any remaining errors.

5 CONCLUSION

This paper identifies ethical risks raised by AI implementations in lift systems and considers how those risks may affect the people who use, build, maintain and govern them. It does not claim to provide finished solutions, but it sets out objectives that future solutions will need to meet if AI-enabled lift systems are to be fair, safe and accountable. It also argues that ethical reflection should become a normal part of technical research in the lift industry, rather than something considered only after systems have already been developed.

AI-enabled lift systems of the future should be fair, privacy-preserving and environmentally responsible. They should be supported by clear accountability structures, practical routes to challenge or exit systems, and governance arrangements that allow those affected by the technology to have meaningful influence over its development and operation. Workers in the lift industry, as well as the passengers who rely on lifts, should not be treated as passive subjects of automation.

As AI permeates the lift industry, a choice must be made. The technology can be allowed to develop without sufficient ethical scaffolding, risking surveillance, inequality, deskilling, responsibility avoidance and eventual enshittification. Alternatively, the industry can build that scaffolding early and deliberately. If it does, AI may help produce lift systems that are safer, more accessible, more sustainable and more accountable than those we have today.

6 REFERENCES

[1] M. Appleby, "Artificial Intelligence in Lift Systems," in *Symposium on Lift and Escalator Technologies*, Kettering, 2025.

[2] S. Wachter, B. Mittelstadt and C. Russell, "Why Fairness Cannot Be Automated: Bridging the Gap Between EU Non-Discrimination Law and AI," in *Computer Law & Security Review*, 2021.

[3] C. Holm, "Smart Buildings: Law and Ethics," in *Digitalization of the Construction Industry*, 2018.

[4] L. You, J. Zhou, Z. Li and F. Chen, "AI Ethics in Smart Homes: Progress, User Requirements and Challenges," arXiv:2412.09813, 2024.

[5] R. Smith, "The Internet of Things, Big Data, Machine Learning, and the Lift & Escalator Industry," *Symposium on Lift & Escalator Technologies,* no. 5, 2015.

[6] L. Bainbridge, "Ironies of automation," *Automatica,* vol. 19, no. 6, pp. 775-779, 1983.

[7] HM Government, "Low Carbon Construction: Innovation and Growth Team – Final Report," Department for Business Innovation and Skills, London, 2010.

[8] E. Bird, J. Fox-Skelly, N. Jenner, R. Larbey, E. Weitkamp and A. Winfield, "The Ethics of Artificial Intelligence: Issues and Initiatives," in *European Parliamentary Research Service*, Brussels, 2020.

[9] International Energy Agency, "Energy and AI," in *IEA*, Paris, 2025.

[10] R. Kenny, "Water Use in AI and Data Centres," Department for Science, Innovation and Technology, London, 2025.

[11] C. Doctorow, "Enshittification: Why Everything Suddenly Got Worse and What to Do About It," Verso Books, London, 2025.

[12] M. Elish, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction," Engaging Science, Technology, and Society, 2019.

[13] OECD, "Recommendation of the Council on Artificial Intelligence," OECD, 2024.

BIOGRAPHICAL DETAILS



Matthew Appleby is a software engineer at Peters Research Ltd. He is involved in project managing the new Elevate Online and leads front-end development. He joined Peters Research in 2019 and graduated with a first-class (Hons) degree in Software Development in 2024.