How housing associations are adapting to AI

NATIONAL HOUSING FEDERATION

Foreword

Artificial intelligence (AI) is rapidly reshaping how housing associations across the country think, plan, and deliver services. Over the past year, many of our members have shown a growing appetite to understand the opportunities AI offers, while also navigating the practical and ethical challenges that come with this new technology.

You told us you're seeking new ways to innovate and improve the services you offer to residents, and it's clear that AI is seen as a great opportunity for this. At the same time, you've also shared the very real challenges that come with adopting AI. This includes concerns around data privacy, navigating governance and securing buy-in, or the current lack of a clear sector-wide strategy in the UK. We recognise these hurdles and understand the uncertainty they can create.

Taking this on board, the NHF and our strategic partner Phoenix Software, an IT services provider, surveyed housing associations across the UK to understand their appetite, capability, and concerns when it comes to AI.

The State of AI in Housing 2025 survey revealed:

- **47%** of members who responded to their survey are using AI in day-to-day operations.
- **30%** are not using AI.
- **23%** are not using AI currently but plan to.

The survey also found that:

- **87%** of members who responded reported having low knowledge of AI.
- **44%** of respondents have no AI policy in place.

This shows there's opportunity for learning and improvement within the sector.

Working alongside NHF members already on their AI journey, we put together this report to share best practice and learnings from across the sector. From data accuracy and supporting earlier interventions, to exploring more efficient internal processes, members are testing AI in ways that are thoughtful, measured and grounded in strong governance. I hope this report is useful for our members who are getting started on their AI journeys.

I would like to say a huge thank you to all the members who contributed. All your experiences are so insightful and will help many housing associations transform the services they offer residents in the years ahead.

Simon Nunn Executive Director of Member Relations at National Housing Federation

Contents

Executive summary

Across the sector, members are considering where AI could be most beneficial to their organisations. We feature seven case studies in this report. They bring to light the opportunities AI presents to housing associations, in increasing efficiency and productivity, while improving resident experience.

In this report, we outline the UK's AI policy background, introducing the relevant legal frameworks and regulatory developments. The UK is seeing rapid growth in AI development and adoption. Various existing legal frameworks affect the deployment and governance of AI in practice. These include intellectual property laws, consumer and competition laws and human rights laws.

Currently, the government have approached AI through a non-statutory framework, based on a set of five core principles for existing regulators, including safety, transparency, fairness, accountability, and redress. The lack of comprehensive regulation and the complexity of AI systems mean that housing associations face increased uncertainty and risk.

Later this year, the government is expected to bring forward a UK AI Bill, establishing centralised, comprehensive and cross-cutting regulation on AI. This would mark a significant shift from the current approach.

AI usage is far outpacing governance and regulatory initiatives, presenting significant ethical concerns that we explore in this report, including:

- **Data privacy and security risks**, particularly around personal and sensitive data.
- **Bias and discrimination**, where AI systems may reinforce inequalities if trained on unrepresentative data.
- **Lack of accountability and transparency,** the root of many AI decisions are not explained or remain obscure.
- **Misinformation risks** stemming from inaccurate outputs.
- **Environmental impact** due to the energy and water demands of AI infrastructure.
- **Impact on housing developments**, as in some areas data centres are blocking developments, due to electricity grid constraints.

The British Standards Institute (BSI) has warned that many UK firms lack robust AI governance despite rising investment, risking operational failures and reputational damage.

Housing associations featured in this report share examples of how good AI governance mechanisms help to mitigate these concerns. The organisations featured also explored the possibility of using AI to improve efficiency and accuracy in areas such as data handling, compliance checks, and resident communication and service provision. Key recommendations drawn from them include:

- **Start small and secure leadership buy-in** to build the right culture around AI.
- **Create strong governance procedures and policies** that balance risk management and compliance with innovation.
- **Engage with staff early on**, gauge awareness and concerns, and tailor communications and training accordingly.
- **Proactively manage risks** around data security, accuracy and bias.

(Please note opinions expressed here by the organisation are not of AI experts)

Introduction

The UK is experiencing a significant rise in Artificial Intelligence (AI) development and adoption. AI technology is evolving globally at a rapid pace, fuelled by advances in computing technology and an increase in data production.

The rise in the UK was spurred on by unprecedented investment and a comprehensive government strategy. So far, more than £78bn was invested in the UK's AI sector. The government is allocating another £4bn over the next four years to implement the recommendations of its AI Opportunities Action Plan, published in January 2025.

This includes establishing a new Sovereign AI unit, setting up AI Growth Zones and scaling up AI infrastructure. The government is also promoting free AI training for all, accelerating AI awareness and usage.

There are more than 5,800 AI companies operating in the UK, following an 85% increase over the past two years. 39% of UK businesses are already using AI in some way, while another 31% are seriously considering it.

With sufficient oversight and governance, AI can help housing associations to streamline their operations and improve resident services. AI systems can unlock new capabilities and support data-driven decision making for strategic planning and resident engagement.

There are also significant challenges and risks in using AI. Currently, AI lacks comprehensive regulation on both a national and sector level, which could expose housing associations to risk. AI is costly to implement and may not always be well received by residents or confidently used by staff members. There are a variety of ethical considerations to consider, including structural biases in AI systems, data privacy and misinformation. The environmental impact of AI also needs to be considered in the context of the sector's ambitious net zero and climate goals.

Purpose of report

AI is a new area of work for the NHF. Our members have identified NHF support on AI as a priority and it now forms part of our five-year business strategy.

In our conversations, you highlighted significant challenges and uncertainty around AI adoption. These included risk management around data usage and privacy, getting governance buy-in and a lack of clear AI regulation or policies in the UK that outlines housing associations' responsibilities when it comes to implementing AI. We also heard you would like to learn more about how others in the sector are approaching AI.

This best practice report draws on the experiences of housing associations, offering practical examples and valuable insights for others in the sector looking to take the next step in their AI journey.

Now we at the NHF are not experts in AI. It is important that organisations consider their own individual risks and responsibilities when reading this report. It serves as inspiration for our members AI journey's rather than a prescriptive instruction manual.

Policy background

Overview

The UK has approached the regulation of AI through a non-statutory framework, based on a set of five core principles. Sector-specific regulators enforce these principles within their respective domains. Aside from this, a patchwork of related legal frameworks contribute to the AI regulatory landscape.

There is no centralised AI regulator or specific overarching AI legislation in force currently, but the government has signalled its intentions to establish a more robust regulatory framework, and a comprehensive new AI Bill is anticipated later this year. While the EU AI Act did provide an international regulatory framework, it is not applicable to the UK following Brexit.

AI Regulation White Paper (March 2023)

The UK's principal AI regulatory framework was established through the AI Regulation White Paper, published in March 2023 by the previous Conservative government. It presents an explicitly pro-innovation approach to AI regulation and outlines five cross-sector principles for existing regulators to interpret and apply within their respective domains.

Sector regulators are expected to ensure that their respective regulations incorporate these principles, and they are encouraged to consult technical ISO standards when doing so, including ISO42001 and ISO22989. These are certified, international specifications for comprehensive AI governance and they contain helpful lists of AI terms and definitions.

The five principles laid out in the paper are:

1. Safety, security and robustness.
2. Appropriate transparency and explainability.
3. Fairness.
4. Accountability and governance.
5. Contestability and redress.

Relevant legal frameworks

In addition to this, various existing legal frameworks affect the deployment and governance of AI in practice. These include intellectual property laws, consumer and competition laws and human rights laws.

Anti-discrimination laws, such as the Equality Act 2010, hold particular significance for AI governance as they provide a legal basis for accountability and the prevention of systemic inequality and unfair treatment through AI usage.

UK data protection laws are also particularly relevant for AI. UK General Data Protection Regulation (GDPR) and the Data Protection Act 2018 govern personal data use, requiring AI systems to be lawful, fair, transparent and accountable.

The Data (Use and Access) Act 2025, which received Royal Assent in June 2025, introduced new provisions and reformed data laws to help foster innovation, especially regarding automated decision-making (ADM) and copyright for AI training.

This more permissive regime has authorised a wider range of circumstances in which personal data can be used within ADM. However, there remain tight safeguards on the use of sensitive personal data (such as race, gender or health). The Information Commissioner's Office (ICO) publishes extensive and helpful guidance on AI and data protection.

Future developments

In October 2025, the UK Department for Science, Innovation and Technology (DSIT) announced plans to establish an AI Growth Lab as a new blueprint for AI regulation.

These labs will serve as regulatory sandboxes, in which firms can trial AI products in real-world conditions, with certain rules and regulations temporarily relaxed under supervision.

The government has outlined sample use cases, including the streamlining of housing development approvals. Successful pilots will inform future regulatory reforms, through guidance updates, codes of practice or statutory changes.

Later this year, the government is expected to bring forward a UK AI Bill, establishing centralised, comprehensive and cross-cutting regulation on AI. This would mark a significant shift from the current approach.

The anticipated Bill is expected to include provisions for safety, accountability, and transparency, with a focus on high-risk systems. Stricter testing requirements, better accountability measures for AI models, and rules for copyright and data usage are likely to feature.

It is also expected to incorporate specific workplace protections, such as rights to explanation and human reconsideration for high-risk decisions, and a ban on certain emotion-recognition technology in employment settings.

Implications for the social housing sector

The lack of comprehensive regulation and the complexity of AI systems mean that housing associations face increased uncertainty and risk.

It is the responsibility of the Regulator of Social Housing to interpret and apply the core principles for the social housing sector, but it is unclear how it has approached this so far. Without intermediate guidance from the Regulator, housing associations must interpret and apply the White Paper directly, which comes with numerous challenges and ambiguities.

There have been calls for the sector to come together to demand better regulatory frameworks that reflect the unique needs and risks of social housing. This might involve more proactive engagement with regulatory bodies to help shape a sector-specific framework and governance model.

Ethical concerns and background

Despite the opportunities that AI offers, it also presents significant ethical risks that must be recognised and effectively mitigated through AI governance mechanisms. These include a robust AI policy, clear guidelines on acceptable use, AI literacy training and AI security awareness training.

Data bias and privacy concerns

Without adequate safeguards, AI has the potential to undermine privacy rights and facilitate data misuse. Poorly secured AI systems can lead to data breaches that expose users to both legal and ethical harm. Many AI-driven tools rely on the mass analysis of personal data, as well as surveillance, monitoring and behavioural prediction. This raises real concern around privacy and security.

AI models can also produce discriminatory outcomes that reinforce existing inequalities. AI systems are trained on human-generated data, which inevitably contain societal biases, such as gender or racial prejudices. The generic datasets that AI systems are trained on may not reflect the diverse communities served by the social housing sector, for example. If not carefully managed, AI can perpetuate and amplify these biases in real-world applications, such as housing allocations.

Accountability

The 'black box' nature of AI technology poses risks when it is entrusted with decision-making. Many AI systems, especially deep learning models, are obscure, with no easy way of understanding how they reach specific outputs. Without appropriate mechanisms for explaining and justifying decisions, errors and bias are more likely to remain undetected.

Misinformation

AI exacerbates the continued challenges of digital misinformation and disinformation, which we've seen over the past decade. AI systems are known to 'hallucinate' and generate outputs that seem confident and plausible but are untrue. This may include factually incorrect, biased or nonsensical information, such as fabricated sources, dates and events.

Presenting inaccurate information as fact has many detrimental consequences, including eroding public trust and weakening the value of truth in our society. It can also put users at direct risk of harm, for example when false information concerns health conditions and medical advice.

AI technology can also be weaponised to build disinformation campaigns, where fake information or content is intentionally and persuasively designed to deceive audiences. In these instances, AI is often employed to distort reality, exacerbate polarisation and seed conflict. AI-driven scams or propaganda campaigns often target the most vulnerable in society.

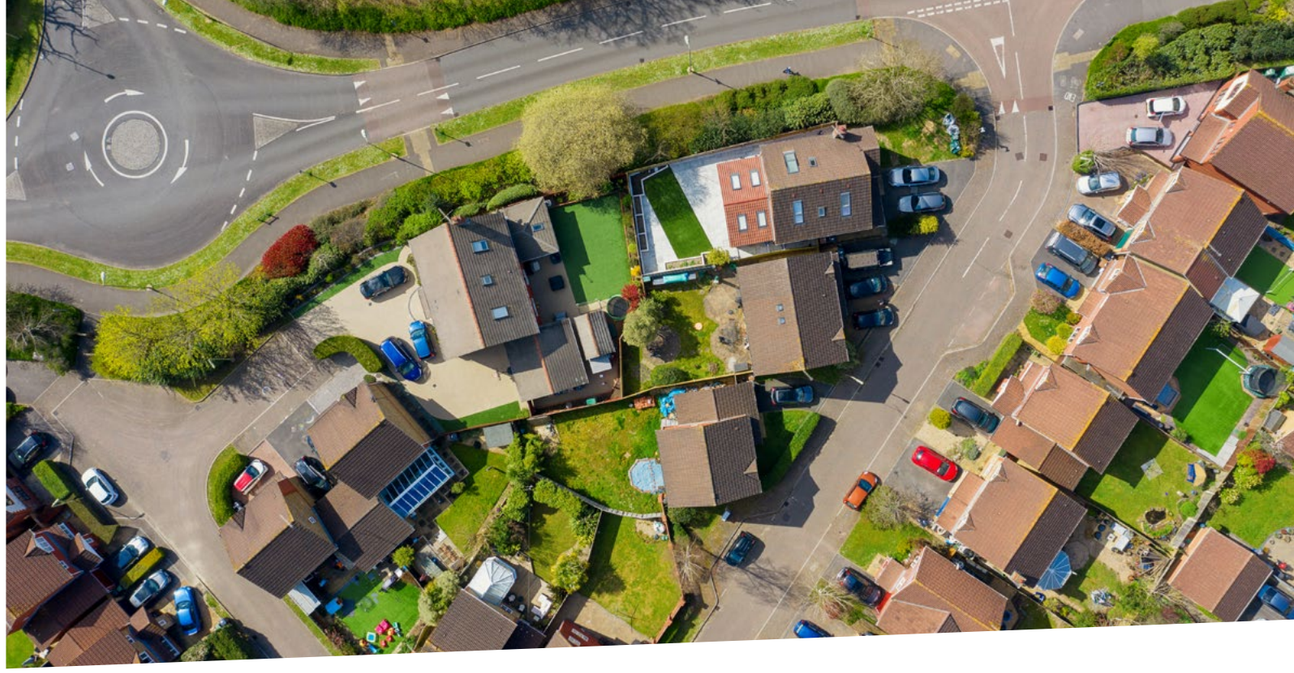

Environmental considerations

AI usage and growth pose significant climate and environmental risks. These risks primarily come from the substantial infrastructure needed to build and operate AI systems, including energy and water-intensive data centres and critical mineral extraction.

To power their complex computing systems, data centres that host AI technology use a lot of energy, which in most places is still generated from the burning of fossil fuels. The International Energy Agency has reported that a request made through ChatGPT consumes 10 times the electricity of a Google search. The agency has also estimated that in the tech hub of Ireland, the rise of AI could see data centres account for nearly 35 per cent of the country's energy use in 2026.

Data centres also require huge volumes of water to cool electrical components. Globally, AI-related infrastructure may soon consume six times more water than Denmark, a country of six million people, according to academic research. As data centres tend to be highly concentrated in certain locations, they can threaten local water supplies, placing pressure on nature and people. Data centres also rely on critical minerals and rare elements, which are often obtained through dirty mining procedures.

In addition to the direct environmental impacts, rapidly increasing demand places AI infrastructure in competition with critical societal and environmental needs, such as the expansion of renewable energy infrastructure needed to mitigate climate change.

Impact on housing development

In some locations, data centres are blocking housing developments, due to electricity grid constraints. For example, energy needed by a series of data centres on the Isle of Dogs, London, has derailed council plans to build more social housing. Tower Hamlets Council has the highest child poverty rate in the whole country, largely due to the lack of affordable and decent housing in the area. The Council admitted that there was a "severe" risk that "housebuilding, at scale, is unable to proceed for potentially 10+ years due to lack of available electricity capacity". In 2023, a study by the Council found as much as 75% of all electricity demand in the area was coming from data centres, versus 16% for offices and 6% for homes.

In 2022, a similar problem was reported in West London, with London City Hall warning developers that a growth in data centres along the M4 corridor would potentially halt housing development for over a decade.

Governance and regulation

AI usage is growing rapidly, far outpacing governance and regulatory initiatives. The British Standards Institute (BSI) has warned that many UK firms lack robust AI governance despite rising investment, risking operational failures and reputational damage. On a national scale, the absence of comprehensive, cross-sectoral regulation means that AI technology may not be receiving sufficient oversight, especially given the complex ethical risks it presents.

Where NHF members and the sector are now

Already, AI is on many housing associations' radar. These organisations are aware of the opportunity AI presents, not only in increasing efficiency and productivity through streamlining operations and getting more done with fewer resources, but also improving the customer experience.

There's a lack of existing AI policies and guidance to protect against pitfalls concerning data security and GDPR. This uncertainty is leading to caution when it comes to AI adoption across the sector and huge variation in where housing associations are on their AI journeys.

A sector-wide survey, The state of AI in housing | Phoenix Software, carried out by Phoenix and the NHF shows that while nearly half of housing associations who replied to the survey are using AI on a daily basis, 87% of respondents rate their AI knowledge as low and 60% have no AI strategy. Meanwhile, 30% of respondents are not using AI and 23% are not using AI currently but are planning to. Demonstrating an overall appetite to embrace AI, this highlights there are opportunities for learning and improvement within the sector, with many already exploring or preparing to implement AI into their work.

In the survey, of those who weren't using AI, 62% of respondents were from the smallest housing associations with fewer than 249 staff. This highlights size of organisation as a key factor impacting AI adoption. In some cases, smaller housing associations may be more limited in the time and resource they can offer towards new projects, contributing towards the variation in AI adoption between housing associations.

What housing associations are aligned on, however, is the importance of understanding what other organisations are doing in this space. This would help housing associations to understand AI and where it could be most beneficial (for both colleagues and tenants), explore different ways of approaching implementation, and in some cases, take the first steps in beginning their AI journey. This report aims to do this by sharing good practice examples from across the sector, looking at how some housing associations have approached AI adoption within their organisations.

Key themes

Housing associations featured in the report were motivated to explore the possibility of AI to improve efficiency and accuracy in areas such as data handling, compliance checks, resident communication, and service provision. Although most of the case studies were still in the early stages of their AI journey, some already recognised positive impacts for staff and residents, such as improved satisfaction, more effective resource allocation and the ability to make more targeted service improvements.

Across the case studies, we can draw out some key learnings for other organisations exploring the opportunities of AI. These are to:

- **Start small and secure leadership buy-in** to build the right culture around AI.
- **Create strong governance procedures and policies** that balance risk management and compliance with innovation.
- **Engage with staff early on**, gauge awareness and concerns, and tailor communications and training accordingly.
- **Proactively manage risks** around data security, accuracy and bias.

Getting started

Many of the case studies talked about starting with a small and simple AI project, while securing buy-in from senior leadership and the board, and getting the necessary governance and policies in place.

The ideal first project is low risk but high impact, such as automating administrative tasks. It is most likely an internal-facing project, to avoid risks to residents. This could also be in the form of a pilot or proof-of-concept to persuade leadership, and staff, of the benefits of AI in a relatively low risk way.

Additionally, Manningham suggest organisations should start with the problem that they need to fix, rather than the tool they want to use. Orwell emphasise the importance of a "culture first" approach, making sure that robust policies, training, and safeguards are in place before rolling out AI.

Governance and policies

Strong governance is seen as a fundamental foundation for good use of AI. This should define the parameters for how the organisation engages with AI, balancing risk management against flexibility for innovation.

This requires the early involvement of relevant personnel across the organisation, not just AI experts, but also IT, data protection, governance, and service leads. For Eastlight, this included the creation of an AI working group comprising of senior leaders and managers. Saxon Weald ensured they had the right expertise at board level, recruiting a board member with extensive experience of transformation in multinational corporations.

The case studies also developed AI policies and undertook compliance reviews and AI impact assessments to maintain robust governance and ensure regulatory and legal compliance. For example, Eastlight worked with data protection consultants to develop a Generative AI policy that balanced innovation with compliance, covering responsible use, data security and ethical considerations.

Staff engagement

Early and continuous staff engagement is a key theme in the case studies. Some sought to gauge staff awareness and apprehension towards the roll out of AI. For example, Eastlight conducted a survey asking about current confidence levels, feelings about current AI usage and ideas for how AI could improve residents' services.

Staff communications and training focused on addressing gaps in knowledge and awareness and alleviating concerns, for example, around job security or "new system risk". Manningham talked about being transparent in communications, specifically around what AI would and wouldn't do, as well as creating safe spaces to share feedback.

But staff engagement was also about creating desire – generating excitement about the opportunities created by AI to improve, rather than replace, their jobs, as well as to improve services for residents. For example, at Incommunities staff use AI as a tool to polish, simplify and humanise complaints' letters, to ensure a consistent tone of voice. In this way AI is used as a tool for staff, rather than as a replacement of staff.

Risks & mitigations

A key concern for housing associations featured within the case studies was around data security, such as the risk of personal or sensitive data being entered into AI systems. Robust governance procedures and policies were used to mitigate these risks, including strict rules on what can and can't be shared and technical safeguards, such as automated redaction. Manningham's AI system is hosted in Canada, creating additional risks which were managed through early due diligence, a compliance review, supplier assurances, and documentation and risk assessments.

Additionally, the case studies raised risks around data accuracy and bias. A key lesson from Orwell's AI staff training programme was that while AI is powerful, it is fallible. It is important to establish robust data governance, to mitigate data accuracy issues as AI is only as good as the data it uses. Recognising this, Together Housing invest time and resource in the data cleaning process to ensure their AI predictive model is based on high-quality and accurate data. For additional assurance of reliability, they have a 'human-in-the loop' to 'reality check' the outputs of the model.

The importance of including a human verification process was also recognised by Your Housing Group who have set up multiple test environments and developed analytical solutions to validate the results, which are then assessed by a human moderator.

A few of the case studies talked about the usefulness of being involved in AI in Housing peer networks with housing providers facing similar challenges. By sharing experience, learning and best practice, the networks helped them to avoid making repeated mistakes.

Case studies

Collated here are seven case studies, from a range of different social landlords. They include large and small housing associations, with homes in different areas of the country.

Saxon Weald

Saxon Weald is a housing association managing approximately 7,000 properties across Sussex and Hampshire. Their mission is to provide high-quality homes and services, with a strong focus on customer experience, proactive service delivery, and operational efficiency.

In their journey to adapting to AI, Saxon Weald focused on gaining senior sponsorship and building strong foundations with key stakeholders. In early 2025, Saxon Weald introduced a formal AI policy to ensure the responsible, ethical, and legally compliant use of AI.

Case study

Inspired by internal and external developments, Saxon Weald decided to get started with their AI journey. They quickly recognised that building strong foundations with key stakeholders played an important role in their digital transformation.

When Debbie Chun, Chief Information Officer (CIO) joined Saxon Weald, she was tasked to review legacy applications and develop a strategy to leverage modern technology and enhance service delivery.

This work was motivated by several factors including increasing regulatory requirements and new legislation, such as Awaab's Law, heightened cyber security risks and financial pressures to "do more with less," particularly in relation to value for money. Driven by these developments, the potential benefits of AI became clear, particularly for automating transactional and repetitive tasks, freeing up time for higher-value activities.

To prepare for AI they started to build strong foundations with key stakeholders. Securing board buy-in and sponsorship became a critical first step. The use of AI carries significant risk and needs to align with a board's risk appetite.

They recruited a new board member with extensive transformation experience in multinational corporations, ensuring the right expertise was represented at the highest level. This led to the launch of their transformation program.

When tackling risks and challenges, the board remained central to resolution.

Debbie Chun notes: **"The biggest challenge we faced was fear, which could undermine confidence in the new system."**

She raised these concerns with their board programme sponsors, who were incredibly supportive. They encouraged her to move forward while ensuring appropriate mitigations were in place.

Once they had board-level buy in, Saxon Weald began introducing the concept with their colleague consultative group and other senior roles.

Understanding that AI could raise concerns about job security, they emphasised that its purpose was to reduce the burden of repetitive, non-value-added tasks, not to replace people.

They were also aware that AI was already being used informally by colleagues, which indicated a strong appetite for innovation.

The next step was creating a formal AI Policy that established clear guardrails for AI use and adoption. For instance, the only approved AI tool was Microsoft Copilot, and it could not be used to make specific decisions about customers.

The policy underwent rigorous review by their infrastructure and security team and GDPR lead, before being presented across the business and approved by the executive team and board.

To complement the policy, they introduced an AI Impact Assessment process. This ensures any requests for AI use outside the current policy are formally evaluated, maintaining robust governance while allowing flexibility for emerging opportunities.

Following its approval, they launched a company-wide campaign to clarify permitted tools and appropriate use of AI. This included live demonstrations, guidance sessions, and communications highlighting both the benefits of AI and the associated risks.

After this, Saxon Weald then integrated AI into their new housing management solution by implementing Brik AI Smart Assistant. This innovation consolidated 18 customer-facing inboxes into one, enabling the AI assistant to read, summarise, and route customer contacts to the correct team while applying appropriate response targets.

This was a significant improvement: previously, over 1,500 emails per quarter were spread across 18 inboxes, with no visibility for their customer services team and highly variable response times.

The accuracy of case routing has been excellent, the customer service team now has full visibility of all customer contacts, and they have reliable monthly data on response times. The first performance report showed that 98% of customer contacts were responded to on time, a significant improvement.

Adoption has been exceptionally smooth. Their teams appreciate that cases are now automatically created, saving them valuable time. They've reached a point where teams are asking for more. AI usage is growing, and there's genuine excitement about its potential.

Eastlight Community Homes

Eastlight Community Homes Limited is a vibrant, resident-led housing association for the East of England. The organisation looks after more than 14,500 homes across a healthy mix of social rent, affordable rent, and shared ownership. Eastlight's core purpose is to provide safe, affordable, and high-quality homes and neighbourhoods, while working in true partnership with residents.

As AI became more widespread, Eastlight considered how it could harness the technology to address residents' evolving needs. Starting their AI journey about a year and a half ago, the organisation's aim was to improve efficiency, by automating routine processes, minimise human error to ensure greater accuracy, and optimise workflows, enabling staff to dedicate more time to meaningful resident engagement.

Case study

Eastlight's approach to introducing AI was structured and collaborative. The first step was securing buy-in from board and senior leadership, which was essential to ensure alignment with their strategic goals. From there, they began exploring how to implement AI to support their organisation responsibly and effectively.

In 2024, Eastlight's board participated in a dedicated session with a specialist consultation service focused on generative AI, which creates new content and ideas, including conversations, stories, images, videos, and music.

It explored Eastlight's strategic appetite for AI, including potential opportunities and risks, and provided insight into how AI could enhance operations and decision-making. Feedback from the session indicated an overall positive appetite for leveraging AI within the organisation, laying the foundations for future innovation and responsible implementation.

Following this session, Eastlight embedded an 'Is AI used in this report?' section into all board and committee papers to ensure transparency and accountability in the adoption of AI.

The housing association launched the 'Think Customer, Think Smarter: AI @ Eastlight' project, an internal AI working group, run by an executive sponsor, comprising of senior leaders and managers that have committed to looking at the ways in which AI can continually enhance residents' services.

Eastlight then worked closely with their dedicated data protection consultants to develop a Generative AI Policy. This involved merging existing policies with new guidance specific to AI. The policy covers key areas including responsible use, data security, and ethical considerations.

On the technology side, the organisation introduced LayerX, a tool designed to manage and monitor AI usage securely. LayerX helped with setting up policies to block certain platforms, redirecting staff to approved tools like Microsoft Copilot, and reinforcing data inputting rules to avoid entering personal, and sensitive business information, into AI systems. Eastlight also implemented automatic redaction for any data that might be entered inadvertently, adding that extra layer of protection.

One of the biggest challenges for the organisation was balancing innovation with compliance. They recognised that a strict, risk-conscious policy could restrict innovation.

To mitigate this, they conducted a survey to gauge colleague awareness, identify concerns, and learn how people might want to use AI in their roles. They assessed confidence levels, if AI could improve services for residents, and how employees felt about current AI usage.

Only 40% of staff felt confident about using AI, with some concerns about misuse, data, and job security. These insights proved invaluable, helping to shape Eastlight's decision to introduce compliant but flexible policies and safeguards.

The housing association then tailored their communications and training to build trust and show the real benefits of AI in everyday work. They raised awareness by sharing best practice for using AI, which included blog posts, posters, all staff emailers, training, and face-to-face sessions. Eastlight also committed to providing more bespoke training to build confidence and ensure staff know how to use AI effectively and safely.

Bringing AI into Eastlight has already delivered tangible benefits. AI has streamlined repetitive tasks, reduced administrative burden, and improved accuracy in areas such as data handling and compliance checks.

Staff shared that using Microsoft Copilot for drafting reports and summarising data has reduced turnaround times. In one example, Eastlight's Home Solutions team (repairs and maintenance service) is exploring and embracing AI as a virtual assistant to improve first-time fix rates, enhance safety, and support faster decision making on repairs.

Beginning with low-risk, high-impact areas like automating admin tasks or summarising reports has built confidence and demonstrated value quickly amongst the team.

While the use of AI is developing quickly, Eastlight remains mindful of both the opportunities and the risks it brings. As a housing association, Eastlight must balance the efficiency AI can offer with the need for resident trust, transparency and clear accountability.

Eastlight is also aware that residents are increasingly using AI tools to help them contact the housing association about repairs, service requests and complaints. Where information is inaccurate or tools are misused, this can require additional resource to respond in a way that meets resident needs and its regulatory obligations. For this reason, Eastlight continues to seek both internal and external insight to manage risk and prepare for future change.

This is an early step in Eastlight's approach to AI adoption, with a clear focus on trust, transparency, efficiency and compliance.

Manningham Housing Association

Manningham Housing Association (MHA) is a community-based housing association in Bradford, providing safe, affordable homes and neighbourhood services, with a strong focus on customer care and inclusion.

MHA has actively explored practical, low-risk AI use cases since 2024. They have a formal AI Policy, provide staff guidance on safe use of generative AI, and run awareness sessions.

Recently, the housing association modernised their telephony system by moving from a legacy phone set-up to a cloud-based contact centre with automation and AI-enabled capabilities.

Case study

MHA's primary goal, in introducing a new telephony system, was to improve their residents' experience. Increased efficiency and better management of data were key additional benefits.

Their previous telephony arrangements limited their ability to manage call demand, provide consistent service at peak times, and generate meaningful performance insights. MHA also sought stronger business continuity and flexibility, so customer services could continue seamlessly during office moves, system outages, or changes to working patterns.

A cloud-based system potentially offered shorter wait times, fewer missed calls, and better call-backs. It could also make the service more resilient. Automation features also enable more intelligent call handling, such as clearer interactive voice response (IVR) options, skills-based routing, call recording for quality, and data dashboards.

MHA approached the implementation of a new telephony service like a change programme, requiring organisation wide change, rather than a pure 'tech upgrade'. Moving to a cloud telephony platform with AI capability would affect how staff work day-to-day, how residents experience contact, and how MHA managed risk and data.

The process began by mapping the customer contact journey to identify friction points and defining success criteria such as faster response and clearer call handling, better reporting, and options to support residents who need extra help.

The cloud transition included resilience planning with fallback routes, rehearsals, and IT collaboration to ensure readiness. MHA worked closely with an external IT team, their own data colleagues, and senior leadership team to ensure the solution supported business continuity, compliance, and value for money.

Along their AI journey, organisational and resident buy-in was key. MHA presented a business case to senior leadership focused on resident outcomes and compliance. An executive sponsor was assigned to the programme.

Staff were apprehensive of the new system. Some feared increased monitoring would reduce call quality and had anxiety regarding replacing human judgement.

MHA alleviated staff concerns about AI and monitoring by being transparent that AI's role would be to support call routing and insight but not make key decisions. They also stressed that they did not intend to use AI sentiment as a sole basis for staff performance decisions.

They shared quick demos highlighting the advantages for their teams, fewer misrouted calls, easier call transfers, clearer queues. To create safe spaces for feedback, they used drop-ins, short surveys, and named change champions.

Throughout the process, MHA reassured frontline staff by reinforcing the team's role in shaping the quality service.

Resident engagement focused on service experience. MHA tested messaging such as IVR wording, ensured accessibility needs were considered, and monitored feedback through existing channels. Where direct resident involvement was limited, the housing association used customer insight and frontline feedback as a proxy and committed to incorporating resident feedback on the service after launch.

Another challenge MHA faced were data residency concerns. The new telephony system would be hosted in Canada. This raised legitimate questions regarding residents' personal data and call recordings. To tackle data risks, the organisation performed a Data Protection Impact Assessment (DPIA) and put appropriate international transfer safeguards in place, including contractual protections and security assurances.

For MHA, strong governance matters. Early involvement from IT, governance, and data protection colleagues helped manage the issue and reassure staff that decisions were being made responsibly. These teams completed a diligence and compliance review and assessed supplier assurances and documentation. They used clear internal governance and sign-off processes, performed a risk assessment and agreed mitigations, and were explicit about what personal data is processed in the telephony platform.

Migrating to the cloud has increased resilience and flexibility at MHA, supporting hybrid teams and future digital integration. The new system offers insights to drive quality, but outputs must be carefully checked by staff due to variable accuracy with regional accents and multilingual callers.

Another initial challenge MHA faced was in-built biases in the telephony sentiment analysis, an AI-powered technology that uses a caller's language and tone to identify their emotional state. Since the system was hosted in Canada, it struggled with UK accents and multilingual customers, creating inaccurate sentiment scoring and reducing trust in the tool's outputs.

To mitigate this, MHA reframed sentiment as an indicator, not a verdict, and validated outputs against real call listening and frontline feedback. They also adjusted how they used reporting, and challenged the vendor on training data, testing methods, and error rates.

Bringing in a new system has had a noticeable impact on both resident experience and staff confidence. Clearer routing and queue logic has helped get residents to the right team more reliably. Better call management features, like transfers, call-backs, and queue visibility, improved the day-to-day experience for frontline staff and reduced manual workarounds.

There's also improved operational visibility. The system provides stronger reporting on call volumes, peak times, and outcomes, which helps managers plan resources and respond quickly to pressure points. Patterns are emerging around repeat contact reasons, bottlenecks, and missed calls, enabling more targeted service improvements.

Overall, the new system modernises telephony while ensuring AI is used as a supportive tool, not a sole decision-maker.

Orwell Housing

Orwell Housing's purpose is to make a positive difference in the communities it serves across East Anglia by providing safe, comfortable homes, alongside the care and support that customers rely on. The organisation manages over 4,000 homes for 8,000 people and delivers specialist care and support to 1500 residents within its services.

Over the past year, Orwell Housing recognised that AI presented substantial advantages and opportunities to enhance operational efficiency and, crucially, improve the delivery of services to its customers. Their journey into AI began in 2024 with an "AI readiness" review conducted by a Microsoft partner. This review established that, although there were no immediate obstacles to adopting AI, the presence of legacy unstructured data and the requirement for enhanced data controls indicated the organisation was not yet prepared for a full-scale AI rollout.

Case study

Following the readiness review, Orwell's approach to AI is deliberately cautious and risk-based. The organisation focused on establishing formal AI policies to define organisational culture and expectations, clearly outlining appropriate and inappropriate AI uses.

Colleagues had previously accessed various AI tools in a controlled "read-only" mode, Orwell's first significant step towards active AI use was enabling Microsoft Copilot for a select group eager to explore the technology. Their experiences offered invaluable insight into practical applications and highlighted risks that required careful management.

One of the greatest challenges was the pace of change driven by Microsoft which necessitated swift action.

Introducing AI was not without risk, especially in relation to maintaining data security. The possibility of personal or sensitive data being entered into AI systems without adequate controls to prevent data loss became clear. Ideally, the organisation would have implemented more technical controls prior to a full rollout.

To mitigate this, the organisation adopted a "culture first" approach, embedding robust policies and training to guide responsible AI use, while continuing to review technical safeguards.

These measures were deemed essential to balance the benefits of AI with the organisation's responsibilities regarding data protection, security, and maintaining customer trust.

Training began with leadership before being rolled out to colleagues via Teams-based drop-in sessions. Each session was guided by the AI policy, which outlined both appropriate and inappropriate uses of AI, with practical examples drawn from personal and professional scenarios to make the guidance accessible. The sessions were well received, helping to build awareness and confidence throughout the organisation.

Delivering training on AI usage and policy across a diverse team with varying levels of experience was always going to be a challenge. While some colleagues had confidently used AI in their personal lives, others were encountering it for the first time. Feedback was overwhelmingly positive, as the sessions highlighted both the opportunities and significant risks associated with AI.

A key lesson from the AI training programme was helping colleagues understand that AI, though powerful, is fallible.

Alongside expanding Microsoft Copilot's use across the organisation, Orwell Housing drew on lessons learnt and sector best practice to develop its AI policy and training plan.

In early 2025, Orwell Housing participated in a sector-wide research project titled "AI in Social Housing, Aspirations and Applications." This project involved conducting colleague surveys and interviews to gauge perceptions of AI, the opportunities it could offer, and the concerns it raised.

Orwell Housing also actively participated in several AI in Housing peer groups too. Since housing providers face similar challenges, these groups played a crucial role in shaping the organisation's approach to AI. The group's collective knowledge helped Orwell Housing avoid repeated mistakes and accelerated decision-making.

Their experience of safe implementation leads them to recommend the benefits of engaging with peers, collaborating openly, and participating in sector groups to accelerate decision-making and strengthen approaches.

Housing providers face unique challenges, such as safeguarding vulnerable customers, managing sensitive data, and balancing technology investment with funding for frontline services.

Collective action ensures that everyone can take full advantage of AI's benefits while managing associated risks responsibly.

Together Housing

Together Housing is one of the largest housing associations in the north of England. They manage nearly 39,000 homes over a wide geographical area and provide a range of support services that affect 80,000 people. The organisation offers quality homes for rent, retirement living, and extra care. They also build homes for sale through shared ownership.

Data plays a crucial role in helping Together Housing to deliver excellent services. Adopting AI has helped them improve the quality of their data, enabling them to make data-driven decisions and develop services that are more efficient, targeted, personalised and proactive.

Case study

A key challenge for Together Housing was maintaining accurate, rich data collection by resident facing teams, which is essential for targeting interventions effectively. They turned to AI modelling to identify residents at risk of losing their homes.

Ending a tenancy is categorised as preventable when it involves issues that the organisation can address, such as abandonment, neighbour disputes, or antisocial behaviour. It becomes non-preventable when it arises from factors beyond a housing association's control, such as changes in residents' circumstances.

For that reason, Together Housing's main objective with the project was to develop a predictive model that can identify tenancies at risk of ending. This would allow the housing association to take preventive measures, reduce tenancy turnover, improve tenant retention, and minimise the costs associated with ending a tenancy.

Together Housing devised a carefully constructed methodology that draws upon concepts and tools from survival analysis, neural networks, and likelihood estimation. They refer to this data modelling approach as Survival Likelihood Deep Learning.

A neural network is a machine learning model inspired by the structure and function of the human brain. It consists of layers of interconnected nodes (or neurons) that process and learn from data. Like the human brain, neural networks are highly flexible and versatile, capable of learning complex patterns and making accurate predictions.

However, the AI models operate in a "black box" closed and opaque environment. It can be challenging to understand how they arrive at their conclusions as their inner workings are not always transparent, making it difficult to pinpoint exactly where or how specific knowledge is stored or applied.

For Together Housing, the neural networks played a crucial role in learning from large datasets to identify patterns and make predictions about residents at risk of losing their homes. The network's ability to handle complex, non-linear relationships in the data allowed Together Housing to achieve high levels of accuracy and flexibility in predicting tenancy outcomes.

Housing data was collected from internal databases, including detailed information on each tenancy, the primary tenant (including payment history), the property, and the household.

The data was then filtered, cleaned, and encoded for subsequent analysis using a custom statistical method based on neural networks.

Operational teams were then able to work with residents at risk. For example, income officers contacted residents with suggestions for alternative methods of paying rent.

To evaluate the model's accuracy, Together Housing "time travelled" into the past, forecasting end of tenancies based on historical data.

Once predictions were made, they compared them to actual outcomes by coming "back to the future" (the present) and assessing whether the predicted end of tenancy occurred as forecasted.

The accuracy score is reported as the percentage of correct predictions, it measures how often the model correctly predicted that a tenancy would fail within a year or continue beyond a year. Based on testing with six-year-old dates, the model is 84% accurate. Measuring preventability is 81% accurate.

Together Housing found that the model proactively reduced the risk of vacancies and associated financial losses. It enabled residents to receive more targeted support, improving their satisfaction and stability. The model also improved resource allocation, ensuring that support is provided where it's most needed.

Since the prediction model has been in place, it's been essential in understanding residents and making every contact count. By getting this right and identifying tenancies at risk, putting in targeted interventions, supporting residents, and preventing a significant number of tenancies from failing, they have saved over £3m the last two years.

Additionally, the model can predict when current active tenancies are likely to terminate overall, which has helped to inform policies and resource planning across the housing association.

Your Housing Group

Your Housing Group (YHG) provides safe, affordable homes to people at every stage of life. They manage more than 30,000 homes across the north of England. Through their passion for housing, more people have a safe place to call home.

For more than five years, YHG has drawn on AI-driven approaches. This began with the use of advanced Structured Query Language (SQL) techniques, powerful methods for querying and analysing large datasets, which allowed the organisation to model and replicate complex decision-making processes.

The organisation has expanded its AI capabilities, over the past year, creating an income housing management analytics solution to support their rent arrears management. This succeeded with the help of AI tools like Microsoft Copilot and Power BI automated insights.

Case study

Prior to exploring automated insights, YHG relied on a rent arrears management tool provided by an income analytics provider. With this contract coming to an end, the housing association began to explore cost-effective alternatives that offered better oversight of the tool. They discovered AI could give them the ability to construct an inhouse solution from the ground up, tailored to their needs.

Before building any model, the organisation first defined what outcomes they wanted, and what rules should drive the decisions. They expressed these requirements in plain business language so that both technical and non-technical audiences could understand and agree on how the model should behave.

Drawing on the knowledge, skills and technology within their income and data teams, YHG began creating a system that would autonomously process arrears data by replicating complex, human decision-making to prioritise cases efficiently. They referred to this system as Income Management Analytics Solution (IMAS).

IMAS retained human oversight for final decisions. This meant it could be adapted quickly and safely to changing external environments, such as the recent transition for residents from Housing Benefit to Universal Credit.

The project timelines were ambitious, with the minimum viable product delivered in under three months. Rather than seeking to simply demonstrate that the technology could work in principle, YHG prioritised an ongoing process of refining and enhancing their solution.

The modelling work was demanding, but rapid development and testing cycles enabled the team to ‘fail fast’ and learn quickly. By using data-driven models to test multiple configurations at the same time, they identified the strongest options and repeated at speed.

Creating multi-layered, flexible SQL to handle complex data, and set the right rules for segmentation and prioritisation, was technically demanding. Luckily, the data team had the skills needed to build and set up these advanced systems.

The resulting system now allows for complex segmentation based on tenancy type, customer circumstances, account balances, and payment patterns. IMAS uses reliable income and payments data to make informed decisions on how caseloads are prioritised, and this has led to improved accuracy and service levels.

There were some risks with this in earlier iterations of IMAS. Initially, it was presenting too many priority cases, creating an unmanageable workload. Using the benchmark from the previous system, they identified the number of cases that should have been flagged and the number of cases that were over presenting as priorities.

YHG mitigated this using multiple test environments, developing analytical solutions to validate the results before being assessed by a human reviewer. This resulted in increased accuracy and reduced workload, enabling them to divert resources where they were most needed. The system now also allows for a full audit that gives a position statement of everyone’s account and where they sit, to ensure no resident is missed.

YHG have run IMAS successfully for six years. IMAS was well received by the organisation. It has generated a wealth of operational insights and performance reporting for management and leadership teams whilst securing vital savings in licensing costs.

Incommunities

Incommunities provides 22,000 homes and services across West Yorkshire. Their aim is to create the best experience for their residents, ensuring they feel heard, respected and supported. To achieve this, they collaborate with residents, communities, and local partners.

The housing association has explored the use of AI tools, such as Copilot, to assist colleagues in crafting empathetic communications that have a warm, clear, and consistent tone of voice.

Case study

Inspired by feedback from their residents, Incommunities started to use AI to improve their complaints communications process. The organisation undertook a complaint's blueprinting exercise with residents familiar with the complaint's procedure. Their residents advised that on occasions they felt some responses were not always empathetic.

Incommunities also acknowledged that the tone of voice was not consistent, and sometimes the letters were not very clear.

Determined to improve resident satisfaction and address the challenges in the writing process, the organisation turned to AI. They hoped that integrating AI into their workflow would give their correspondences a consistent tone of voice, reflecting their values of empathy and clarity. Additionally, they felt that this approach would streamline the drafting process, making it smoother and more efficient for colleagues, reducing the time and effort required to write the complaints letters.

To begin developing their communications and customer voice, internal teams at Incommunities worked collaboratively to develop clear guidance documents for letter writing, conducted tone of voice training sessions, and designed a letter writing framework that fostered empathy and clarity whilst reflecting the organisation's brand style and values.

They embedded prompts and rules into Microsoft Copilot to standardise structure and tone. The

team has also created structured templates for complaint letters, ensuring each one captured key elements such as complaint definition, findings, outcomes, next steps, and information about the Housing Ombudsman.

The communications team worked with the customer voice team to bring the new approach to life. They had already developed a tone of voice training framework and rolled it out to the organisation, but many colleagues struggled to break old habits and adapt to the new style. To mitigate this, the new AI framework supported the expectations for the tone of voice and style. Alongside this, the communications team collaborated to develop guides and prompts.

Colleagues now draft letters and use Microsoft Copilot to check that their tone is in line with Incommunities' warm, clear and respectful voice, leveraging built-in settings that maintain consistency and reduce the risk of human error.

This new framework is in its early days and colleagues have embraced the change. One colleague said:

"I have been using the new complaints template and AI guidance for some time now. Initially, it was quite different from what I was used to. I found that it has significantly reduced the time needed to draft letters. It has also created greater consistency in both style and tone across my correspondence."

Incommunities remains aware of the risks associated with integrating AI into their communication workflow, such as inaccurate details or inappropriate tone. To mitigate this, they have put in robust quality assurance measures.

The framework emphasises the fact that letters should still be checked by colleagues. The complaints team also review 80% of letters and carefully check the final drafts.

Incommunities continuously update training, and the complaints team carry out coaching on a 1-2-1 basis. Regular feedback from teams informs ongoing improvements.

Through this framework, the organisation hopes to ensure that every complaint letter is accurate and of a high standard and staff feel supported with the new process.

Looking to the future

By starting small and simple, housing associations can begin to build a positive culture around the use of AI in their organisation and secure buy-in from colleagues at all levels.

A useful starting point is considering the problem you would like to fix or the service you would like to improve, rather than thinking about the tool you would like to introduce.

In turn, this can have the benefit of improving efficiency and streamlining productivity within your organisations, while also transforming the services you offer your residents.

Whether you are a housing association that is new to AI, looking to get started or seeking advice on how to further implement it into your organisation, we hope this report helps you feel more informed on the main policies and ethical considerations to be aware of and on the different ways housing associations can use AI.

Phoenix Software, one of the NHF's strategic partners, can also provide you with thorough advice, training, and support as you progress through your AI journey.

The National Housing Federation will continue to work closely with Phoenix and other bodies, to facilitate knowledge sharing and collaboration to support housing associations on their AI journeys.

With thanks

NHF advisory group:

We would like to thank all NHF members who have fed into the development of this work.

In particular:

AI working group
Carolina Padovezi (Chair) and
all members of the group

Case study organisations

Saxon Weald

Eastlight Community Homes

Manningham Housing Association

Orwell Housing

Together Housing

Your Housing Group

Incommunities

Phoenix Software

Dave Lingwood

Greg Dean

NHF Adopting AI team:

Helen Collins

Kevin Garvey

Cara Green

Lizzi Hearn

Katie Miller

Joanne Muigua

Simon Nunn

Charlotte Rogers

Chloë Rutter

Nadia Tsoukala-Steggell

NATIONAL
HOUSING
FEDERATION