

A design theory for democratising artificial intelligence-assisted indoor air diagnostics through low-code platforms



indooraircartoon.com/2026/07/27/a-design-theory-for-democratising-artificial-intelligence-assisted-indoor-air-diagnostics-through-low-code-platforms

27 July 2026

Indoor Air Cartoon Journal, July 2026, Volume 9, #180

[Cite as: Fadeyi MO (2026). A design theory for democratising artificial intelligence-assisted indoor air diagnostics through low-code platforms. Indoor Air Cartoon Journal, July 2026, Volume 9, #180.]

A DESIGN THEORY FOR DEMOCRATISING ARTIFICIAL INTELLIGENCE-ASSISTED INDOOR AIR DIAGNOSTICS THROUGH LOW-CODE PLATFORMS



Prof., what is low-code AI, and how is it relevant to indoor air practice?

The low-code artificial intelligence (AI) workflow is an externalisation of the indoor air professional's diagnostic reasoning into an executable workflow for supporting value-oriented indoor air problem diagnosis and solving to achieve healthy living through healthy indoor air. It represents how the professional systematically investigates an indoor air problem requiring structured diagnostic reasoning execution.

The workflow defines the questions to ask, the information to collect, the variables to evaluate, the connections and interactions among those variables, the decision rules to apply, the probable causes to investigate, and the recommendations required to improve indoor air quality in ways that support healthy living. Using the low-code platform, each reasoning step is visually connected to its appropriate information sources and expected outputs, enabling executable diagnostic reasoning consistently.



The central premise of the design theory is that Low-Code AI can be intentionally designed to empower non-AI professionals to develop and exercise the cognitive governance required to create, govern, and continuously refine AI-assisted workflows. It does so by stimulating mental model development, encouraging meaningful questioning, enhancing cognitive abilities, supporting sound judgement and informed decision-making, and ultimately guiding actions for value-oriented indoor air problem diagnosis and solving. Current indoor air practice lacks such a framework. The proposed cognitive governance-oriented low-code design theory enabled transparent, adaptable, and scalable AI-assisted diagnostic workflows, improving diagnostic quality and value delivery.

Information may be entered manually during site inspections or automatically retrieved from indoor air sensors, Building Management Systems, Internet of Things (IoT) devices, maintenance databases, weather services, applicable indoor air standards, previous investigation reports, occupants' health records and exposure history, and other connected digital information sources. These inputs drive the executable workflow consistently throughout each investigation.

Hmm! So, Low-Code AI is like ready-to-eat food from 7-Eleven shop. You choose the food (pre-built software capabilities), combine it into a meal (AI-assisted workflow), adjust it to your needs, and serve the finished meal (AI-assisted application solving the problem). That is democratisation.

AI consistently executes the expert-defined diagnostic reasoning. It interprets, contextualises, and evaluates incoming information against predefined relationships and decision rules. It retrieves additional information when required, identifies subsequent diagnostic actions, and guides the investigation pathway as new evidence becomes available. Throughout the process, the indoor air professional retains governance over every diagnostic decision and workflow refinement through transparent review, validation, approval, and continuous improvement for future investigations and healthy living.



Yes! Making AI accessible to and usable for non-AI experts.

REFERENCES

- © Moshood O. Fadeyi 2026 | Contact at: mofadeyi@gmail.com | www.indooraircartoon.com
 Pinho, D., Aguiar, A., and Amaral, V. (2023). *Journal of computer languages*, 74, 101185.
 Rao, N., et. al. (2024). In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (pp. 837-852).
 Sá, D., et. al. (2025). In *EPIA Conference on Artificial Intelligence* (pp. 55-67). Cham: Springer Nature Switzerland.

Note: Software enables a computer or digital device to perform functions to achieve an intended purpose. To do this, the software provides the computer or digital device with programming instructions (code) that tell it what to do.

As software systems became larger and more complex, software engineers developed software development tools to reduce the growing effort required to design, develop, test, debug, deploy, integrate, and govern increasingly complex software systems. These tools include integrated development environments (IDEs), compilers, debuggers, testing tools,

version control tools, computer-aided software engineering (CASE) tools, and, more recently, Low-Code platforms. Like other software development tools, Low-Code platforms are themselves software developed by software engineers. However, unlike conventional software development tools, Low-Code platforms are specifically designed to automatically generate much of the programming code required to develop software for these increasingly complex systems.

Although a Low-Code platform automatically generates much of the underlying programming code, its purpose is not simply to develop software but to enable the creation of a software solution that solves a particular problem. A software solution is more than just programming code. It is a digital solution designed to solve a particular problem. It consists of several interconnected components, including the user interface, workflows, business rules, data model, AI capabilities, and integration with other systems. Examples of other systems include indoor air sensors, Building Management Systems (BMS), Internet of Things (IoT) devices, maintenance databases, weather services, applicable indoor air standards, previous investigation reports, occupants' health records and exposure history (where appropriate), and other connected digital information sources. The underlying software architecture organises these components into a functioning software application.

Low-Code platforms enable domain experts in other fields with little or no software programming expertise to develop these software solutions. They do so by designing diagnostic workflows and selecting, combining, and configuring the pre-built software and AI capabilities needed to support those workflows. The Low-Code platform automatically generates much of the underlying programming code required to create the software application. As a result, domain experts can devote more effort to higher-value cognitive activities. These activities include understanding the problem, designing the solution, exercising judgement, making informed decisions, and governing AI-assisted applications.

Although domain experts do not need to engineer the underlying software components, they must understand the problem to be solved, the workflow or process required to solve it, the inputs and outputs at each stage, the ready-made software components with appropriate AI capabilities that best support each stage, how to configure and connect these components into a coherent workflow, and how to evaluate whether the workflow is producing valuable outcomes.

Professionally cooked ready-to-eat food analogy

7-Eleven food analogy

Low-Code AI equivalent

Ready-to-eat food items available in the store.	Pre-built software components
Customers choose the food items with a clear purpose in mind.	Domain experts in other fields select the required software capabilities with a clear purpose in mind.
Customers decide how to combine them into a meal to achieve a purpose.	Domain experts in other fields design the AI-assisted workflow to achieve a purpose.
Customers adjust the sauces and portions to suit the problem context and the purpose to be achieved.	Domain experts in other fields configure the workflow to suit the problem context and the purpose to be achieved.
A completed meal that fulfils the intended purpose.	AI-assisted application that solves the problem and fulfils the intended purpose.

Low-Code AI is like ready-to-eat food from a 7-Eleven store. Customers choose the food (pre-built software capabilities), combine it into a meal (AI-assisted workflow), adjust it to their needs, and serve the finished meal (AI-assisted application). That is democratisation, that is, making AI accessible to and usable by domain experts in other fields with little or no expertise in AI or software engineering for value-oriented problem diagnosis and solving.

Just as a customer at 7-Eleven does not need to grow rice, raise chickens, or cook every ingredient from scratch to prepare a meal, domain experts in other fields do not need to write computer code from scratch to develop an AI-assisted application. The value lies in exercising **appropriate cognitive governance** to select, combine, and configure the right software and AI capabilities to solve the problem at hand.

Beyond supporting the execution of diagnostic workflows, the built-in AI capabilities developed by AI experts continuously support domain experts in other fields in evaluating workflow quality and contextual appropriateness throughout each investigation. These capabilities analyse investigation outcomes, historical diagnostic records, scientific evidence, applicable standards, and the specific context to identify missing variables, conflicting decision rules, incomplete reasoning pathways, and opportunities for continuous improvement.

In essence, Low-Code AI democratises software development, while cognitive governance democratises high-quality problem diagnosis and solving. The former reduces the mechanistic effort required to develop AI-assisted applications, whereas the latter ensures that the software is intentionally designed, appropriately configured, and continually governed to deliver value-oriented solutions within the specific context of the problem.

.....
Fictional Case Story (Audio – available online)– Part 1 (Preface, Ch 1 & Ch 2)

Fictional Case Story (Audio – available online) – Part 2 (Ch 3)

Fictional Case Story (Audio – available online) – Part 3 (Ch 3 cont'd)

Fictional Case Story (Audio – available online) – Part 4 (Ch 4)

Fictional Case Story (Audio – available online) – Part 5 (Ch 4 cont'd)

Fictional Case Story (Audio – available online) – Part 6 (Ch 5 & Ch 6)

..... Preface

In a world increasingly transformed by Artificial Intelligence, automation, and digital technologies, societies became more connected, information became more accessible, and intelligent systems became progressively more capable. Organisations celebrated technological innovation, researchers generated unprecedented volumes of knowledge, and professionals embraced intelligent tools to improve efficiency, productivity, and decision-making. Yet one fundamental question remained insufficiently explored: as the capability to generate solutions advanced, was equal attention devoted to developing an appropriate understanding of the problems requiring those solutions?

Across engineering, scientific research, applied research, education, government, and professional practice, increasing emphasis on efficiency, productivity, and performance gradually shifted attention towards generating solutions. Comparatively less attention was devoted to developing an appropriate understanding of the problems requiring those solutions. This tendency quietly became embedded within professional practice, shaping how problems were interpreted, evidence was evaluated, and decisions were made. Consequently, technological advancement often outpaced the cognitive discipline required to govern its responsible application. Artificial Intelligence could process vast quantities of information, but it could not determine whether a problem had been understood appropriately, whether assumptions required re-examination, or whether proposed solutions genuinely served human values.

Recognising this growing imbalance within herself, a young engineering student embarked upon a lifelong journey through engineering practice, applied research strengthened with scientific reasoning, education, leadership, and personal transformation. In seeking better ways to solve complex problems, she discovered that the greatest challenge was not developing more intelligent technologies, but cultivating a philosophy of professional practice in which an appropriate mental model governs judgement, judgement governs decision, and decision governs action. Her journey forms the subject of this fiction story.

My name is Janet Koroma. I was born in Freecity, the capital of Leony, a developing country. For many years, I believed that the greatest threat to good judgement was ignorance. I assumed that people made poor decisions because they did not have enough knowledge, education, experience, or technical competence. This belief shaped the way I approached school, work, relationships, and professional development. Whenever something went wrong, I concluded that I needed to learn more. I attended another course, read another book, pursued another qualification, or asked someone whom I considered more knowledgeable. It took me many years to discover that knowledge alone does not guarantee sound judgement.

A person may possess considerable knowledge and still misunderstand the situation before him or her, leading to poor judgement and decisions. That person may collect evidence but interpret the evidence poorly. That person may recognise several possible consequences but focus only on the one that supports an early conclusion. That person may make a decision quickly and construct a justification afterwards. That person was me.

The underlying flaw was the absence of a systematic way of developing and continuously refining my understanding of a situation, interpreting evidence, evaluating potential consequences, and making informed judgements and decisions. At the time, I did not know this flaw existed. Even when its consequences became visible, I blamed circumstances, other people, insufficient information, or bad luck. I was not being dishonest. I genuinely believed my explanations. What I lacked was the ability to step outside my own thinking, examine how I had arrived at my conclusions, and determine whether my understanding, interpretation, judgement, and decisions were actually justified.

My flaw did not begin in adulthood. Its roots can be traced back to the circumstances surrounding my birth and the way I learnt to survive uncertainty as a child. My mother was seven months pregnant with me when she and my father fled a conflict-affected region in search of safety. They did not leave because they had carefully planned a new life. They left because staying had become more dangerous than leaving behind everything they had ever known.

My father worked as a maintenance technician in a small manufacturing plant, while my mother taught mathematics at a local secondary school. They lived a modest but stable life. They had a small home, relatives nearby, secure employment, and hopes for the child my mother was carrying. Within a few months, everything changed. The factory where my father worked closed. Armed groups began controlling the roads, making travel increasingly dangerous. Several families disappeared from the neighbourhood without explanation. One evening, an explosion damaged buildings near my mother's school. The following morning, my father looked at my mother and quietly said that they could no longer wait for conditions to improve. It was time to leave.

They travelled with one suitcase, a folder containing their certificates, a small amount of money, and the constant fear that every checkpoint could bring their journey to an end. My mother later told me that she hardly slept during those weeks. She worried that the stress might harm me. She also feared that I might be born before they reached safety. Whenever she felt me move, she gently placed both hands on her stomach and whispered that I needed to remain patient. She told me we were travelling to a safer place, although she herself did not know where that place would be.

I was born shortly after they arrived in Phyland, a wealthy country with four distinct seasons. My parents were physically safe, but safety did not immediately bring stability. My father's technical qualifications were not fully recognised, and he accepted temporary work wherever he could find it. My mother cleaned offices in the evenings while studying to meet the requirements needed to teach again. We lived in a small rented flat above a busy road. The windows did not close properly, mould repeatedly appeared in the bedroom, and heating the flat consumed a painful share of my parents' income. Years later, I recognised the irony. My earliest home suffered from several indoor environmental problems, yet none of us had the knowledge or resources to understand them properly.

My parents protected me from most of their difficulties, but children often notice what adults do not say. I noticed the unopened envelopes on the table. I noticed my father's silence whenever temporary work ended. I noticed my mother carefully calculating prices before quietly returning items to the shelf. Although I did not fully understand what these moments meant, I came to believe that adults often had to make important decisions without knowing whether they would have enough work, enough money, or what tomorrow would bring. I learnt that life could change suddenly and that hesitation could be dangerous. Decisions often had to be made quickly, even when the available information was incomplete.

This became one of my earliest mental rules: when uncertainty appeared, decide quickly and keep moving. The rule helped me in some situations. I became decisive, independent, and difficult to discourage. Teachers described me as confident. Friends often asked me to make choices for the group because I rarely appeared confused. However, what looked like confidence was often a defence against uncertainty. I disliked remaining in situations where I did not know the answer. Uncertainty reminded me of the instability I had experienced as a child. Consequently, I learnt to reduce uncertainty quickly, often by accepting the first explanation that appeared reasonable.

My father reinforced this tendency without intending to do so. He valued action. If a tap leaked, he repaired it. If a door would not close, he adjusted the hinges. If I complained about a school problem, he asked what I intended to do about it. He rarely had the luxury of prolonged reflection. His life had taught him that action created options, whereas waiting often made circumstances worse. I admired him deeply and adopted his practicality. However, I copied the visible action without appreciating the careful observations and

thinking that preceded his decisions. I saw him acting, but I did not always notice how he first developed his understanding of the situation.

At school, my tendency to reach conclusions quickly was rewarded more often than it was corrected. I performed well in examinations. I memorised efficiently, recognised familiar patterns, and produced answers within the required time. When a question resembled something I had encountered before, I usually performed well. However, when a problem required me to question my assumptions, reconsider my initial interpretation, or develop a deeper understanding of the situation, I sometimes struggled. Yet good grades concealed this weakness. Because my answers were often correct, I assumed my thinking process was equally sound.

The weakness became more apparent in my personal relationships. When a friend failed to return my call, I concluded that she was avoiding me. When my mother questioned one of my decisions, I interpreted it as a lack of trust. When a teacher gave critical feedback, I assumed that the teacher had already formed a negative opinion of me. I often reacted to the explanation I had constructed rather than to the situation itself. When additional information later emerged, I sometimes realised that my judgement had been unfair. I apologised when necessary, but I treated each incident as an isolated mistake instead of recognising a recurring pattern. I still did not realise that the real problem was not the decisions I made, but how I had arrived at them.

By the time I entered university to study architectural engineering, I believed that education would remove uncertainty from my life. The discipline attracted me because it appeared to combine scientific principles with practical solutions. Structural behaviour could be analysed. Indoor environmental performance could be measured. Heat transfer, airflow, daylight, and energy performance could be modelled. Building materials had measurable properties, and engineering systems followed established principles. I imagined that technical competence would enable me to replace uncertain judgement with objective knowledge.

I worked hard and graduated with an upper-second-class honours degree. I was proud of the result, particularly because I was the first person in my immediate family to complete a university degree in Phyland. My parents attended the graduation ceremony wearing clothes they had purchased especially for the occasion. My father took more photographs than anyone else in the hall. My mother cried when my name was announced. At that moment, the hardships surrounding my early life seemed to have produced something worthwhile.

Shortly after graduation, I joined an engineering consultancy specialising in building performance, ventilation, and indoor air quality. I began as a graduate architectural engineer and was initially assigned structured tasks. I reviewed drawings, organised monitoring data, assisted with site investigations, and prepared sections of technical

reports. I enjoyed the work because it connected buildings with the people who occupied them. A ventilation system was not merely ductwork, fans, and control logic. Its performance influenced whether people could breathe comfortably, concentrate effectively, recover in hospital, learn in school, or feel safe in their homes.

My technical confidence grew quickly. I learnt to use monitoring instruments, interpret environmental measurements, evaluate ventilation systems, and compare findings with standards and professional guidance. Senior colleagues praised my enthusiasm and willingness to take responsibility. I rarely waited to be told what to do. Whenever a problem arose, I was usually ready with a solution.

Yet the flaw that had followed me from childhood also entered my professional practice. During one of my earliest investigations, employees in a recently refurbished office reported headaches, throat irritation, and difficulty concentrating. Initial monitoring identified elevated carbon dioxide concentrations during afternoon occupancy. I quickly concluded that inadequate outdoor air supply was the primary cause. My investigation focused almost entirely on the ventilation system. I recommended increasing outdoor airflow, adjusting operating schedules, and reviewing occupancy density.

My recommendations were technically reasonable, but they did not solve the problem. A more detailed investigation led by a senior colleague revealed that several factors were interacting. New furniture was emitting volatile organic compounds. Cleaning products had recently changed. Solar heat gain was increasing afternoon discomfort. Some employees had recently been relocated and were seated close to printers. The ventilation problem was genuine, but it represented only one part of a much larger picture.

My senior colleague did not embarrass me. He simply asked, “Janet, what made you decide that carbon dioxide explained the entire complaint?” I replied by describing the measurements and the relevant ventilation guidance. He listened patiently and then asked, “What evidence would have shown you that your first explanation was incomplete?” I did not have a good answer. He did not answer the question for me. Instead, he smiled gently and said, “Think about it.” Then we returned to the investigation.

For several days, the question stayed with me. However, instead of examining my thinking, I protected myself by reframing the incident. I told myself that junior engineers could not reasonably be expected to identify every contributing factor. I concluded that I simply needed more experience. That explanation was partly true, but it prevented me from recognising the deeper problem. More experience could provide more knowledge, but knowledge alone would not develop a disciplined way of understanding situations, interpreting evidence, exercising judgement, and making decisions.

I continued to develop professionally. After several years, I enrolled in a part-time Master of Science programme while remaining with the same company. The arrangement was demanding. I worked during the day, attended evening classes, completed assignments at

weekends, and sometimes conducted site visits before lectures. The programme expanded my knowledge of building science, environmental assessment, ventilation, exposure, health risk, and research methods. I became better at analysing complex technical problems. My company promoted me, and clients increasingly trusted me to lead investigations.

Completing the MSc should have increased my confidence in my professional judgement. Instead, it created a more subtle form of confidence. I knew more, so I trusted my conclusions more readily. However, greater knowledge also made it easier to defend my initial judgement rather than question it. I could identify evidence supporting my preferred explanation and present technically persuasive arguments. I had become more knowledgeable and technically capable, but I had not yet learnt how to govern the way I developed understanding, interpreted evidence, exercised judgement, and made decisions. As a result, the consequences of my underlying flaw became increasingly serious.

In one hospital investigation, staff reported odours and respiratory discomfort in a treatment area. Monitoring results showed acceptable average pollutant concentrations, and the ventilation system appeared to comply with the applicable design requirements. I initially judged that the reported symptoms were unlikely to arise from indoor environmental conditions. Based on that judgement, I recommended continued monitoring and improved communication with staff.

A nurse later contacted the facilities team and described a pattern that had not been included in the initial complaint records. Symptoms occurred mainly during specific cleaning activities and affected staff with asthma more severely. Short-duration exposure peaks had been masked by the use of average measurements. The underlying cause of the reported complaints could not be explained by the average pollutant concentrations and ventilation performance alone. Instead, it resulted from the interaction between cleaning chemicals, work practices, exposure timing, and the biological vulnerability of the affected occupants.

I felt ashamed when I reviewed the new evidence. My conclusion had not been careless in the ordinary sense. I had followed familiar procedures, reviewed the available measurements, and considered the relevant compliance requirements. Yet I had not developed a sufficiently complete understanding of the situation before exercising judgement. I had treated the absence of evidence in the information I possessed as evidence that the problem did not exist. I had failed to ask what evidence might be missing, whose experiences had not been heard, or how the timing of exposure might influence the interpretation of the available evidence.

More fundamentally, I had not asked myself the question that would later transform both my professional practice and my life: "What is it that I do not know or have not considered

that makes what I know or have considered contextually inappropriate?" For the first time, I could no longer dismiss the pattern as a consequence of insufficient experience. I had knowledge. I had qualifications. I had nearly a decade of professional practice. Yet something fundamental was still missing.

A close friend gradually stopped inviting me to social gatherings and became less responsive to my messages. I concluded that she was offended because I had repeatedly declined earlier invitations while completing my MSc. Instead of asking what had changed in her life, I became defensive. During one conversation, I accused her of making me feel guilty for pursuing my career.

She became quiet before saying, "Janet, I have been going for medical tests. I did not want to worry you until I knew more." Months later, she told me that she had also been caring for her father after he was diagnosed with cancer. The shame I felt was different from professional embarrassment. I had interpreted her behaviour through my own assumptions, fears, and workload. I had not asked meaningful questions. I had not considered alternative explanations. I had judged her before understanding her situation. Fortunately, her medical condition proved manageable, but the incident changed something in me.

I began to recognise the same pattern in different parts of my life. I formed an early explanation, interpreted subsequent evidence through that explanation, underestimated uncertainty, and acted before my understanding was sufficiently mature. When the consequences were uncomfortable, I blamed missing information. Yet I rarely asked whether I had made enough effort to discover what information was missing. Reality had finally confronted me in a way I could no longer ignore.

For several months, I struggled with what to do. Part of me wanted to dismiss the problem as a personality trait. Perhaps I was simply decisive. Perhaps every professional occasionally exercised poor judgement. Perhaps slowing down would make me hesitant and ineffective. My childhood had taught me that delay could be costly, and my career had rewarded my willingness to act.

Another part of me gradually recognised that decisiveness without sufficient understanding could also create harm. The answer was not to avoid judgement or postpone decisions indefinitely. The answer was to develop a systematic way of building understanding before exercising judgement, and to continue refining that understanding as new evidence emerged.

I began experimenting with my own thinking. Before reaching a conclusion, I wrote down what I believed was happening, what evidence supported that belief, what assumptions I was making, what information might still be missing, what alternative explanations remained plausible, and what consequences might follow if I were wrong. I deliberately asked whose perspective I had not yet considered. I distinguished between what I knew,

what I inferred, and what I merely suspected. After making a decision, I reviewed the outcome and reflected on whether the new evidence should refine my understanding, improve my judgement, and influence my future decisions.

At first, this process felt slow and unnatural. I worried that colleagues would interpret my questions as uncertainty. Instead, several colleagues responded positively. Meetings became more productive because assumptions were made visible. Junior staff became more willing to challenge proposed explanations. Clients understood why additional evidence was necessary. My decisions did not become perfect, but they became more transparent and defensible. Most importantly, I became more willing to change my judgement without feeling that changing my mind represented failure.

I gradually realised that the capability I was developing extended beyond individual cognitive abilities such as critical thinking, reflective thinking, abstract reasoning, logical deduction, and creative imagination. I was learning how to govern the way those abilities were used. I began by developing and continuously refining a mental model of the situation.

That understanding, i.e., mental model developed, stimulated meaningful questions. The questions exercised my cognitive abilities and exposed weaknesses in my mental model. The refined mental model then supported better judgement. Judgement informed decisions. Decisions guided action. The outcomes and lessons from those actions, in turn, refined my mental model. Although I did not yet call this process cognitive governance, that was what I was developing. At the time, I simply knew that I was becoming less afraid of uncertainty because I had learnt how to work with it rather than merely escape from it.

..... Chapter 2

After twelve years in the company and the indoor air industry, I had become a senior professional with responsibility for complex investigations, staff development, and aspects of digital innovation. By then, my personal journey of learning to develop understanding before exercising judgement had fundamentally changed the way I approached professional problems. It also changed what I noticed.

Around the same time, the company had begun investing in Artificial Intelligence-assisted diagnostic tools. These systems could organise information, classify complaints, compare measurements with thresholds, and generate recommendations. They appeared powerful, but I increasingly recognised that they supported information processing far better than they supported the governance of professional reasoning. That distinction enabled me to recognise limitations that I used to overlook.

When scientific evidence changed, the diagnostic logic did not change automatically. Indoor air professionals could explain what needed to be updated, but software developers had to translate that reasoning into code. The process required meetings, specifications,

development schedules, testing, approval, and deployment. By the time the updated diagnostic logic became operational, professional understanding might already have moved on.

I also noticed that the systems often captured conclusions more effectively than the reasoning that produced them. They could record that a ventilation intervention had been recommended, but they did not always make visible how competing explanations had been evaluated. They could compare a pollutant concentration with a threshold, but they did not necessarily represent occupant vulnerability, exposure duration, building context, uncertainty, or the consequences of alternative actions. Valuable professional reasoning remained within the experience of individual practitioners, meeting notes, or conversations that disappeared after the investigation. The problem felt strangely familiar.

The organisation possessed scientific knowledge, technical expertise, software systems, and experienced professionals. Yet it lacked a systematic way for those professionals to directly represent, govern, adapt, and continuously refine their reasoning within AI-assisted diagnostic workflows. The organisation's limitation mirrored my own. Just as I had once moved from incomplete understanding to judgement without a disciplined governance process, AI-assisted systems could move from available data to recommendations without adequately representing how professional understanding should develop and evolve.

I began discussing the issue with colleagues. Some believed that more advanced algorithms would solve it. Others argued that the company needed more software developers. Both solutions could help, but neither addressed the question that continued to trouble me: Why were the professionals responsible for indoor air diagnosis unable to directly govern how their own reasoning was represented, adapted, and continuously refined within the AI-assisted systems designed to support them? After all, they were the ones who generated that reasoning, applied it in practice, and continuously refined it as new scientific evidence and professional experience emerged.

The more I explored Low-Code AI, the more I realised that the technology addressed only part of the problem. It could enable indoor air professionals to participate directly in building and modifying AI-assisted diagnostic workflows, but it could not ensure that the reasoning embedded within those workflows was scientifically sound. A poorly governed mental model could still be transformed into a poorly governed workflow. The real need, therefore, was not simply to make AI development easier. It was to develop a governance-oriented design theory that guided how professional knowledge and reasoning should be represented, questioned, refined, validated, governed, and continuously improved through Low-Code AI.

The idea gradually evolved into a proposal for an industrial Doctor of Engineering. My company recognised that the problem I had identified affected both the quality of indoor air investigations and the long-term value of its AI-assisted diagnostic systems. If qualified

indoor air professionals could directly govern how their professional reasoning was represented, adapted, and continuously refined within AI-assisted workflows, the organisation could respond more rapidly to emerging scientific evidence, preserve valuable professional knowledge, strengthen transparency, and deliver greater value to clients and building occupants. These potential benefits aligned closely with the company's long-term strategy for digital innovation.

After several discussions with senior management, the company agreed to co-sponsor the doctorate in partnership with the university. The research would address a genuine industrial problem while contributing new engineering knowledge with practical value beyond a single organisation. To support the project, my workload was adjusted where possible, selected responsibilities were redistributed within the team, and sufficient flexibility was provided for data collection, university activities, and research supervision while I continued working full-time. The arrangement allowed the research to remain firmly grounded in authentic professional practice.

For me, the decision required deeper reflection. I had already completed a part-time MSc while working full-time and understood the personal commitment another degree would demand. More importantly, I realised that the industrial problem I wanted to investigate reflected a flaw that had shaped my own life. The research was no longer simply about improving AI-assisted indoor air diagnostics. It had also become an opportunity to understand and further overcome the way I developed understanding, interpreted evidence, exercised judgement, and made decisions.

With the support of both my company and the university, I prepared my Doctor of Engineering application. The research would investigate a real industrial problem while contributing a governance-oriented design theory for democratising AI-assisted indoor air diagnostics through Low-Code AI. The research problem that formed the basis of my Doctor of Engineering application was as follows.

“Indoor air quality (IAQ) problem-solving aims to improve healthy living by enabling healthy indoor air through scientifically informed engineering decisions. Achieving this objective requires qualified indoor air professionals to integrate multidisciplinary evidence and interpret complex building conditions. They must then exercise sound professional judgement, make informed decisions, and develop value-oriented interventions that effectively address the underlying causes of indoor environmental problems.

As scientific knowledge, building technologies, occupant expectations, and regulatory requirements continue to evolve, AI-assisted diagnostic systems have increasingly been proposed as decision-support tools to enhance professional practice. Despite their potential, there remains a noticeable gap between the capabilities expected of AI-assisted indoor air diagnostics and their performance in real-world engineering practice.

The practical problem addressed by this research is the gap between the current performance of AI-assisted indoor air diagnostic systems and the desired performance required to support adaptable, transparent, scientifically rigorous, and value-oriented indoor air problem-solving. Although advances in Artificial Intelligence have improved computational capability, many AI-assisted diagnostic systems remain difficult to adapt as scientific understanding of the indoor environment under investigation evolves through new evidence relating to indoor air quality, building performance, occupant health, and diagnostic practice.

Updating diagnostic logic frequently requires specialist software development. This creates delays between new scientific evidence and its implementation within operational AI systems. Consequently, qualified indoor air professionals often remain dependent on software developers to modify AI-assisted diagnostic workflows, even though they possess the expertise required to determine how those workflows should evolve.

This dependency creates several interconnected barriers. Scientific knowledge continues to develop through new research findings, revised standards, changing regulations, completed investigations, and accumulated professional experience. However, AI-assisted diagnostic systems may not evolve at the same pace. Workflow modifications often depend on software engineering processes rather than professional judgement. As a result, AI-assisted workflows may temporarily fail to reflect current scientific understanding. This limits their adaptability and long-term usefulness.

Additional concerns relate to transparency, organisational learning, and governance. AI-assisted recommendations may not always make the underlying professional reasoning visible to engineers, organisations, regulators, or building owners. When professional reasoning is insufficiently represented, organisations may struggle to preserve valuable professional knowledge and explain diagnostic decisions.

This also limits multidisciplinary collaboration, knowledge transfer to less experienced professionals, and the continuous improvement of AI-assisted diagnostic workflows. Consequently, professional knowledge may remain embedded within the experience of individual practitioners instead of becoming reusable organisational knowledge that supports continuous improvement.

These challenges suggest that the principal limitation may not lie solely in the computational capability of Artificial Intelligence. Instead, it may arise from how professional knowledge and reasoning are represented, governed, adapted, and continuously refined within AI-assisted diagnostic systems. If this is the case, improving algorithms alone may be insufficient to close the gap between current and desired performance. A different AI development paradigm may therefore be required.

Low-Code AI platforms provide a potential opportunity to address this problem. They may enable qualified indoor air professionals to participate directly in developing, adapting, and

governing AI-assisted diagnostic workflows without requiring specialist software programming expertise. However, several important questions remain unanswered.

It is not yet clear whether the scientific, technical, organisational, and cognitive limitations of current practice are sufficiently significant to justify such a paradigm. Neither is it clear how AI-assisted diagnostic workflows should be designed and governed by qualified indoor air professionals while maintaining scientific rigour, transparency, explainability, accountability, and value-oriented reasoning.

Furthermore, limited empirical evidence exists to determine whether AI systems developed directly by qualified indoor air professionals using Low-Code AI can outperform conventional software-engineered AI systems in terms of adaptability, transparency, scalability, governance, diagnostic quality, and value delivery.

Accordingly, there is a need to investigate the limitations of contemporary indoor air problem-solving practice. There is also a need to develop a governance-oriented design theory that enables qualified indoor air professionals to directly design, adapt, and govern AI-assisted diagnostic workflows using Low-Code AI.

Finally, there is a need to empirically evaluate whether this approach provides a superior paradigm for AI-assisted indoor air diagnostics compared with conventional software-engineered AI systems. Addressing this problem has the potential to establish a new governance paradigm for AI-assisted decision support in indoor air diagnostics and other knowledge-intensive professions where expert reasoning must continuously evolve alongside scientific knowledge.”

This need forms the basis for the research questions and hypotheses that guided my D.Eng. study.

(i) To what extent do contemporary indoor air quality (IAQ) problem-solving practices exhibit scientific, technical, organisational, and cognitive limitations that justify the development of a new governance-oriented design theory enabling qualified indoor air professionals to directly develop, adapt, and govern AI-assisted diagnostic systems using low-code platforms?

(ii) How should qualified indoor air professionals directly design, adapt, and govern AI-assisted diagnostic workflows using Low-Code AI while maintaining scientific rigour, transparency, explainability, and value-oriented reasoning?

(iii) To what extent does low-code AI automation developed directly by indoor air professionals improve the adaptability, transparency, scalability, governance, diagnostic quality, and value delivery of indoor air problem-solving compared with conventional software-engineered AI systems?

For the first research question, the Null Hypothesis (H_{01}) is that the scientific, technical, organisational, and cognitive limitations of existing indoor air problem-solving practices do not justify a design theory for democratising Artificial Intelligence-assisted indoor air diagnostics through low-code platforms.

The Alternative Hypothesis (H_{11}) is that the scientific, technical, organisational, and cognitive limitations of existing indoor air problem-solving practices significantly justify a design theory for democratising Artificial Intelligence-assisted indoor air diagnostics through low-code platforms.

For the second research question, the Null Hypothesis (H_{02}) is that the proposed governance-oriented design theory does not significantly enable qualified indoor air professionals to directly design, adapt, govern, and continuously refine AI-assisted diagnostic workflows using Low-Code AI while maintaining scientific rigour, transparency, explainability, and value-oriented reasoning.

The Alternative Hypothesis (H_{12}) is that the proposed governance-oriented design theory significantly enables qualified indoor air professionals to directly design, adapt, govern, and continuously refine AI-assisted diagnostic workflows using Low-Code AI while maintaining scientific rigour, transparency, explainability, and value-oriented reasoning.

For the third research question, the Null Hypothesis (H_{03}) is that the proposed design theory for democratising Artificial Intelligence-assisted indoor air diagnostics through low-code platforms does not significantly improve adaptability, transparency, governance, scalability, diagnostic quality, or value delivery compared with conventional software-engineered AI systems.

The Alternative Hypothesis (H_{13}) is that the proposed design theory for democratising Artificial Intelligence-assisted indoor air diagnostics through low-code platforms significantly improves adaptability, transparency, governance, scalability, diagnostic quality, and value delivery compared with conventional software-engineered AI systems.

The research questions and problems informed the following objectives of my D.Eng. research:

- (i) To investigate the scientific, technical, organisational, and cognitive limitations of contemporary indoor air quality (IAQ) problem-solving practices in order to determine whether they justify the development of a governance-oriented design theory that enables qualified indoor air professionals to directly develop, adapt, and govern AI-assisted diagnostic systems using Low-Code AI.
- (ii) To develop a governance-oriented design theory that enables qualified indoor air professionals to directly design, adapt, and govern AI-assisted diagnostic workflows using

Low-Code AI while maintaining scientific rigour, transparency, explainability, and value-oriented reasoning.

(iii) To evaluate the extent to which Low-Code AI automation developed directly by qualified indoor air professionals improves the adaptability, transparency, scalability, governance, diagnostic quality, and value delivery of indoor air problem-solving compared with conventional software-engineered AI systems.

..... Chapter 3

Research Methods

Methods for Research Question 1:

Overview

Research Question 1 examined whether the current ways of solving indoor air quality (IAQ) problems have important scientific, technical, organisational, and human-thinking limitations that justify developing a new design theory. This design theory proposes that qualified indoor air professionals, rather than software developers, should be able to directly create and improve AI-assisted diagnostic systems using low-code platforms that require little or no computer programming.

As the purpose of this study was to determine whether such a new theory was needed, rather than to develop or test an AI system, a sequential mixed-methods research design was adopted. The study first collected numerical evidence through an international survey and then gathered deeper insights through interviews with experienced professionals and detailed analyses of real indoor air diagnostic cases. Combining these different sources of evidence allowed the findings to be verified from multiple perspectives, ensuring that the conclusions reflected real-world professional practice rather than relying only on published literature.

The study was conducted over a period of twelve months. Data were collected from professionals working in the fictional countries of Aurelia, Norlandia, Pacifica, and Westoria. These countries were intentionally created to represent different levels of technological maturity in indoor air quality practice. Each country was assumed to have well-established building engineering industries, mature indoor air quality regulations, active professional communities, and widespread use of digital technologies for assessing building performance. Studying professionals from these different countries also helped determine whether the challenges identified were common across the profession rather than being unique to a particular country or organisation.

Study Settings

The study was carried out in 32 operational buildings where indoor air quality (IAQ) problems are commonly investigated as part of normal building operation, maintenance, commissioning, and occupant complaint management. The buildings comprised four commercial office buildings, four hospitals, four educational institutions, four hotels, four shopping centres, four laboratories, four residential apartment buildings, and four industrial facilities. Using real buildings instead of laboratory experiments ensured that the study reflected the practical challenges that indoor air professionals face in their daily work. The selected buildings represent the types of environments where professional IAQ investigations are most commonly undertaken. Each building type presents different indoor air challenges. For example, hospitals require strict control of airborne contaminants to protect patients and healthcare workers.

Office buildings primarily focus on maintaining indoor environments that support occupant comfort and productivity. Educational institutions experience changing occupancy levels throughout the day, creating varying ventilation demands. Residential buildings are concerned with the long-term exposure of occupants to indoor pollutants because people spend extended periods of time in their homes. Industrial facilities often face additional challenges from pollutants generated by manufacturing processes and other industrial activities. Including different building types allowed the study to examine a broad range of indoor air problems under realistic operating conditions.

Participants were recruited from 24 organisations, comprising engineering consulting firms, facility management companies, environmental consultancy organisations, government agencies, healthcare engineering departments, universities, and research institutions. These organisations represent the key groups responsible for investigating indoor air problems, operating and maintaining buildings, developing technical guidance, enforcing regulations, and advancing indoor air research. Recruiting participants from different sectors ensured that the study captured a wide range of professional experiences, organisational practices, and approaches to diagnosing indoor air problems.

Conducting the study across 32 buildings and 24 organisations provided a realistic representation of contemporary IAQ practice while ensuring diversity in building functions, ventilation systems, occupancy characteristics, pollutant sources, and organisational decision-making processes. This broad range of settings increased confidence that the findings reflected common challenges faced by indoor air professionals rather than the practices of a single organisation or a particular building type. Consequently, the study settings provided a robust and realistic foundation for determining whether a new design theory is needed to enable qualified indoor air professionals to directly develop, adapt, and govern AI-assisted diagnostic systems through low-code platforms without relying heavily on software developers.

Participant Recruitment

Participants were experienced professionals who regularly investigated and solved indoor air quality (IAQ) problems as part of their daily work. Their roles included identifying the causes of indoor air problems, collecting and interpreting air quality measurements, assessing building ventilation systems, recommending corrective actions, evaluating whether remediation efforts had been successful, and ensuring that buildings met relevant indoor environmental standards. Including professionals with different responsibilities provided a broad understanding of how IAQ problems are identified and managed in practice.

To ensure that participants had sufficient practical experience, only professionals with at least five years of relevant work experience and involvement in a minimum of twenty completed IAQ investigations during their careers were included in the study. These requirements helped ensure that participants had encountered a wide range of indoor air problems and had developed sound professional judgement through real-world practice rather than relying mainly on academic knowledge or limited field experience.

Participants were recruited through professional organisations, national indoor air societies, facility management associations, engineering institutions, and environmental health networks. These organisations were selected because they bring together professionals actively working in indoor air quality, building engineering, environmental health, and facility management. Invitations to participate were distributed electronically through organisational mailing lists and professional social media platforms, allowing experienced professionals from different organisations and regions to participate in the study.

A total of 612 professionals initially responded to the invitation. After assessing whether they met the eligibility requirements, 538 professionals were included in the final study, representing an overall response rate of 87.9%. The large number of participants provided a broad range of professional opinions and experiences, increasing confidence that the findings reflected common challenges faced by the indoor air profession.

The final group included mechanical engineers, environmental engineers, industrial hygienists, occupational hygienists, indoor air consultants, building scientists, healthcare engineers, facilities managers, and IAQ researchers. This diversity ensured that the study captured perspectives from professionals involved in different stages of indoor air investigation, diagnosis, and problem-solving.

Quantitative Data Collection

Quantitative data were collected using a structured questionnaire developed specifically for this investigation. A questionnaire was selected because it allowed information to be collected from a large number of experienced professionals in a consistent and systematic manner. Every participant answered the same questions in the same order, making it

possible to compare responses objectively across different professions, organisations, and countries.

Instrument development was informed by an extensive review of indoor air diagnostic literature, expert systems, explainable artificial intelligence, knowledge engineering, decision-support systems, cognitive engineering, organisational learning, and low-code software development. More than 300 peer-reviewed journal articles, conference papers, technical standards, and industry publications published over the previous fifteen years were reviewed to identify the key challenges reported in existing indoor air diagnostic practices.

The questionnaire consisted of 64 questions organised into four main categories, each corresponding to one aspect of the research question. Each category contained 16 questions designed to measure a specific type of limitation experienced by indoor air professionals. Although the questionnaire contained 64 questions, each question was intentionally designed to be concise and directly related to situations commonly encountered by indoor air professionals. This enabled most participants to complete the questionnaire within approximately 25 minutes without requiring extensive written responses.

The scientific limitation category examined issues such as dealing with uncertain evidence, combining knowledge from different sources, understanding how AI reached its conclusions, interpreting scientific evidence correctly, and updating diagnostic methods as new scientific knowledge became available. The technical limitation category examined dependence on software developers, the ability to modify AI systems, flexibility of diagnostic workflows, compatibility with other digital systems, transparency of AI decisions, and the difficulty of implementing changes.

The organisational limitation category focused on how knowledge is shared within organisations, how quickly organisations respond to new information, the effort required to maintain AI systems, collaboration among different professionals, and the ability to expand AI systems across multiple projects or organisations. The cognitive limitation category explored how professionals think through diagnostic problems, develop possible explanations, interpret evidence, make decisions, and convert their professional knowledge into AI-assisted diagnostic processes.

Responses were recorded using a seven-point rating scale, ranging from 1 (strongly disagree) to 7 (strongly agree). A seven-point scale was chosen because it provides sufficient detail to capture differences in professional opinions while remaining simple and familiar for participants to use. Higher scores indicated stronger agreement that a particular limitation existed in current indoor air problem-solving practice.

Before the questionnaire was distributed to the main study participants, it was pilot-tested with 32 experienced indoor air professionals, who were not included in the final study. The

pilot study was conducted to ensure that the questions were easy to understand, free from ambiguity, and relevant to professional practice. Participants were also asked to provide feedback on the wording of the questions, the overall length of the questionnaire, and whether any important issues had been overlooked.

The pilot testing evaluated clarity, readability, internal consistency, completion time, and construct validity. The results showed that most participants completed the questionnaire within 20 to 30 minutes, with an average completion time of approximately 25 minutes. Based on participant feedback and the pilot analysis, minor wording changes were made to improve readability and ensure that the questions could be understood consistently by professionals from different backgrounds.

The questions asked participants to reflect on their own professional experience when investigating indoor air problems. Rather than asking whether they supported the use of Artificial Intelligence, the questionnaire focused on the challenges they currently encounter during everyday practice. For example, participants were asked whether updating AI systems required assistance from software developers. They were also asked whether they could easily modify diagnostic procedures when new scientific knowledge became available.

Another question examined whether they understood how AI systems reached their recommendations. Participants were also asked whether important diagnostic knowledge was lost when experienced staff left their organisations. Additional questions examined whether existing AI systems could be adapted to different building types and changing regulations. Finally, participants were asked whether current diagnostic processes adequately supported professional judgement and decision-making.

These questions were designed to examine four broad areas: scientific limitations, technical limitations, organisational limitations, and cognitive limitations. The purpose of these questions was not to evaluate the participants themselves, but to evaluate the current way indoor air problems are solved. If many experienced professionals reported facing the same difficulties, this would indicate that the current approach has important weaknesses. For example, many professionals might report that they depend on software developers to update AI systems. Others might report that they are unable to understand how AI systems reach their recommendations.

Some may also find it difficult to incorporate new scientific knowledge into existing AI systems. If these challenges are commonly experienced across the profession, they would provide evidence that current indoor air problem-solving practices have important scientific, technical, organisational, or cognitive weaknesses. Identifying these common weaknesses would provide evidence that a different approach is needed. In this study, that different approach is the proposed design theory, which enables qualified indoor air professionals to

directly create, adapt, and continuously improve AI-assisted diagnostic systems through low-code platforms without relying heavily on software developers.

The questionnaire was administered electronically using a secure online survey platform. Electronic distribution enabled professionals from different organisations and countries to participate conveniently without geographical restrictions. The online platform also ensured that responses were recorded automatically, reduced the risk of data entry errors, prevented unanswered mandatory questions, and maintained participant confidentiality throughout the data collection process. The survey platform also prevented duplicate submissions by restricting one completed questionnaire per participant. Participants required approximately 25 minutes to complete the questionnaire. Completed questionnaires were stored in an encrypted database accessible only to the research team for subsequent statistical analysis.

Participation in the study was entirely voluntary. To encourage participation, respondents were informed that they would receive an executive summary of the research findings after completion of the study. This report highlighted emerging trends, common challenges, and recommended practices identified across the participating organisations. Several collaborating professional organisations also endorsed the study and encouraged their members to participate because the findings were expected to contribute to improving future indoor air diagnostic practice. No financial incentives were provided.

Qualitative Data Collection

To complement the quantitative findings, semi-structured interviews were undertaken with 40 internationally recognised indoor air quality (IAQ) experts purposively selected from academia, government laboratories, engineering consultancies, healthcare engineering departments, and multinational building service organisations.

A purposive sampling approach was adopted because the study sought detailed insights from professionals with extensive experience rather than opinions from a large number of respondents. These experts were selected based on their recognised contributions to indoor air research or professional practice, including leadership roles in major IAQ investigations, scientific publications, development of technical guidelines, or extensive experience managing complex indoor air problems. Selecting experts from different sectors ensured that the study captured a broad range of perspectives from research, regulation, engineering design, building operation, and healthcare practice.

Each interview lasted between 60 and 90 minutes and was conducted either face-to-face or through secure video conferencing, depending on the participant's geographical location and availability. A semi-structured interview format was chosen because it allowed every participant to answer the same core questions while giving them the flexibility to explain their experiences in greater detail and introduce important issues that had not been

anticipated by the researchers. This approach enabled the collection of richer and more detailed information than could be obtained from the questionnaire alone.

Interviews explored participants' experiences with current diagnostic workflows, AI implementation challenges, software development dependency, knowledge transfer processes, organisational barriers, cognitive difficulties encountered during diagnosis, and perceptions regarding the feasibility of enabling domain experts to directly construct AI-assisted diagnostic workflows using low-code platforms. For example, participants were asked to describe situations where they wanted to update an AI system but could not do so without assistance from software developers, where valuable diagnostic knowledge was lost when experienced staff left the organisation, or where existing AI systems could not easily adapt to new scientific evidence, changing regulations, or different building types.

They were also asked how these challenges affected the quality, speed, transparency, and reliability of indoor air investigations, and whether allowing qualified indoor air professionals to directly develop and modify AI-assisted diagnostic workflows could address these problems.

Interviews were digitally recorded with participant consent and transcribed verbatim. Recording the interviews ensured that participants' responses were captured accurately without relying on handwritten notes or researcher memory. Verbatim transcription preserved every spoken response, allowing the data to be analysed systematically while reducing the risk of important information being overlooked. Participants were given the opportunity to review their interview transcripts to confirm that their views had been accurately represented before analysis commenced.

Data collection continued until theoretical saturation was achieved, whereby no substantially new concepts emerged from successive interviews. In practice, this meant that the same themes and challenges were repeatedly identified by different participants, indicating that additional interviews were unlikely to generate meaningful new insights. Continuing interviews until this point increased confidence that the study had captured the major scientific, technical, organisational, and cognitive issues experienced by indoor air professionals in current practice.

Diagnostic Case Analysis

To provide additional evidence that was independent of the survey and interviews, a documentary analysis was conducted. Documentary analysis is a research method in which existing documents are carefully examined to identify useful information, recurring patterns, and evidence relevant to the research question. Instead of asking professionals what they remembered or believed, this method allowed the researchers to examine records of actual indoor air investigations that had already been completed in professional practice.

A total of 100 completed indoor air quality (IAQ) investigation reports were obtained from collaborating engineering consultancies and facility management organisations. The reports represented investigations carried out over the previous five years in commercial office buildings, hospitals, educational institutions, hotels, shopping centres, laboratories, residential apartment buildings, and industrial facilities across the participating countries. Each investigation report described how an indoor air problem had been investigated from the initial complaint through to the final diagnosis and recommended corrective actions. The reports typically included information such as the reported indoor air problem, inspection findings, air quality measurements, laboratory test results, ventilation assessments, diagnostic reasoning, recommendations, and the final conclusions reached by the investigation team.

Before analysis, all reports were anonymised by removing confidential information such as organisation names, client details, building locations, project identifiers, and personal information. This ensured that the analysis focused on the diagnostic process while protecting the privacy and commercial confidentiality of the participating organisations. Studying these reports enabled the researchers to observe how indoor air problems were actually diagnosed in professional practice rather than relying solely on participants' opinions or recollections.

Each report was systematically analysed to identify diagnostic workflow characteristics, number of investigative iterations, decision pathways, expert consultation requirements, software tools employed, documentation practices, uncertainty management strategies, and instances where diagnostic reasoning required modification because of emerging evidence. For example, the number of times an investigation strategy was revised before reaching a final diagnosis was examined. It was also examined whether additional experts were consulted during the investigation. The use of AI and other digital tools to support diagnostic decision-making was also examined. In addition, the management of uncertainty when the available evidence was incomplete or conflicting was evaluated.

Finally, it was examined whether new scientific knowledge or evidence required changes to the diagnostic approach before the investigation was completed. The researchers also examined whether important diagnostic decisions were clearly documented so that another professional could understand how the final conclusions had been reached. The objective of this analysis was to determine whether existing workflows demonstrated recurring dependence on specialist software developers, fragmented knowledge translation processes, or limited adaptability when new diagnostic knowledge became available.

In other words, the analysis was undertaken to determine whether the same problems identified through the questionnaire and interviews could also be observed in actual indoor air investigation reports. If the reports consistently showed that professionals depended on software developers to modify AI systems, this would provide evidence that such dependence was a genuine limitation of current practice.

Similar evidence would be obtained if professionals were found to struggle to incorporate new scientific knowledge into existing diagnostic workflows. Evidence would also be provided if AI systems were found to be difficult to adapt to new situations or changing requirements. If these challenges were observed repeatedly across different investigation reports, they would indicate that the identified limitations were widespread throughout current indoor air problem-solving practice rather than being isolated personal experiences or opinions.

Data Analysis

Quantitative analyses were performed using IBM SPSS Statistics Version 30 and AMOS Version 30. IBM SPSS was used to organise and analyse the questionnaire responses, while AMOS was used to examine the relationships between different variables and determine whether the proposed research model was supported by the data. These software packages are widely used because they provide reliable statistical procedures for analysing large datasets. Descriptive statistics summarised participant characteristics and overall response distributions, including participants' professional backgrounds, years of experience, and responses to each survey question. Internal consistency was evaluated using Cronbach's alpha and composite reliability, with acceptable reliability defined as values greater than 0.70.

These tests determined whether questions within each category consistently measured the same underlying concept. Construct validity was assessed through confirmatory factor analysis using maximum likelihood estimation to determine whether the questionnaire accurately measured the four intended categories: scientific, technical, organisational, and cognitive limitations. Model adequacy was evaluated using multiple goodness-of-fit indices, including χ^2/df , Comparative Fit Index (CFI), Tucker-Lewis Index (TLI), Root Mean Square Error of Approximation (RMSEA), and Standardised Root Mean Square Residual (SRMR). Together, these indicators assessed how well the proposed research model represented the observed survey data.

Average Variance Extracted (AVE) values exceeding 0.50, together with the Fornell-Larcker criterion and heterotrait-monotrait (HTMT) ratios, were used to confirm that each category measured a distinct aspect of indoor air problem-solving practice. Structural equation modelling (SEM) was then employed to examine whether the four categories collectively justified the proposed design theory. Statistical significance was determined at $\alpha = 0.05$, indicating less than a 5% probability that significant findings occurred by chance.

Qualitative interview transcripts and diagnostic reports were analysed using reflexive thematic analysis to identify recurring ideas, experiences, and challenges across different participants and investigation reports. Two independent researchers initially coded all transcripts before discussing coding differences to achieve analytical consensus, thereby reducing potential bias. Emergent themes were continuously compared across interviews

and documentary evidence using constant comparative analysis. NVivo Version 15 was used to organise, code, retrieve, and categorise the qualitative data while maintaining a clear record of the analytical process.

Finally, quantitative findings, qualitative themes, and documentary evidence were integrated through methodological triangulation. This involved comparing evidence obtained from the questionnaire, expert interviews, and investigation reports to determine whether they supported the same conclusions. Agreement across the three methods increased confidence in the findings, while any differences were carefully examined before drawing final conclusions regarding whether the identified limitations justified the proposed design theory for democratising AI-assisted indoor air diagnostics through low-code platforms.

Ethical Considerations and Methodology Contribution to Knowledge

Ethical approval was obtained from the host university's Institutional Review Board before the commencement of the study. Participation was entirely voluntary, and written informed consent was obtained from all participants before any data were collected. Participants were provided with an information sheet explaining the purpose of the research, the nature of their involvement, the expected duration of participation, and how the information they provided would be used. They were also informed that they could decline to answer any question or withdraw from the study at any stage without penalty or adverse consequences. No financial incentives were offered for participation, thereby ensuring that participation was based solely on professional willingness to contribute to the advancement of knowledge.

To protect participants' privacy and confidentiality, all survey responses, interview transcripts, and indoor air investigation reports were anonymised before analysis. Personal identifiers, organisation names, building locations, project details, and commercially sensitive information were removed or replaced with anonymous codes. Interview recordings and electronic data were encrypted and securely stored on password-protected institutional servers accessible only to authorised members of the research team. All data were managed, retained, and disposed of in accordance with the host university's research ethics requirements and data governance policies. Findings were reported only in aggregated form to ensure that no participant, organisation, or building could be identified from the published results.

Beyond ensuring ethical compliance, the methodology was intentionally designed to generate original knowledge through the systematic integration of multiple sources of evidence. Quantitative survey data identified the extent to which scientific, technical, organisational, and cognitive limitations were experienced across the indoor air profession. Qualitative interviews provided detailed explanations of why these limitations occurred and how they affected professional practice.

Documentary analysis of completed indoor air investigation reports provided independent evidence of how these challenges manifested during actual indoor air investigations. Collectively, these complementary methods enabled the research findings to be examined from different perspectives while reducing dependence on a single source of evidence.

The principal methodological contribution of this study lies in its integration of quantitative, qualitative, and documentary evidence to establish a robust and reproducible basis for evaluating the scientific need for a new design theory. Rather than relying solely on professional opinions or theoretical arguments, the methodology systematically combined statistical evidence, expert experiences, and documented professional practice.

These different sources of evidence were used to determine whether the identified limitations were genuine and widely experienced across the indoor air profession. They were also used to assess whether the limitations were sufficiently significant to justify democratising Artificial Intelligence-assisted indoor air diagnostics through low-code platforms. This integrated methodological approach provides a transparent framework that can be replicated, adapted, and applied to investigate similar AI-assisted decision-support challenges in other knowledge-intensive engineering and professional domains.

Methods for Research Question 2:

Overview

Research Question 2 sought to determine how indoor air professionals should use low-code Artificial Intelligence (AI) platforms to design, adapt, and govern AI-assisted diagnostic workflows while maintaining scientific rigour, transparency, explainability, and value-oriented reasoning. To answer this question, the study developed and evaluated a governance-oriented design framework that enabled qualified indoor air professionals, rather than software developers, to translate their professional knowledge, experience, and diagnostic reasoning into AI-assisted workflows that could be executed by a low-code AI platform.

Building upon the empirical findings established in Research Question 1, the study transformed the scientific, technical, organisational, and cognitive limitations previously identified into practical design requirements for the proposed framework. In other words, the problems identified in Research Question 1 became the foundation for designing a solution in Research Question 2.

Consequently, the study adopted a Design Science Research (DSR) methodology integrated with research-through-design and iterative expert validation. This approach was selected because the objective was not to evaluate an existing software application. Instead, the objective was to systematically design, refine, and validate a governance-oriented framework. The framework was intended to enable indoor air professionals to

develop, modify, and govern AI-assisted diagnostic workflows without requiring software programming expertise.

The investigation was conducted over an eighteen-month period. Conceptual development and iterative refinement of the governance-oriented framework were undertaken at the host university. Practical implementation and validation were subsequently undertaken in collaboration with engineering consultancies, healthcare organisations, facility management companies, and environmental engineering organisations.

Authentic indoor air diagnostic scenarios were developed from completed indoor air investigations conducted in operational buildings. These scenarios were then used to implement and evaluate the proposed governance-oriented framework under realistic professional conditions. Working across both academic and professional environments ensured that the framework was scientifically rigorous while remaining applicable to real engineering practice.

Development of the Design Framework

The governance-oriented design framework was developed through four iterative design cycles. The first cycle translated the findings of Research Question 1 into the initial design requirements for the framework. The second cycle developed the first working version of the framework and its associated AI-assisted diagnostic workflows. The third cycle implemented and evaluated the framework with experienced indoor air professionals using realistic indoor air diagnostic scenarios. The fourth cycle refined and validated the framework based on the findings obtained during implementation and evaluation.

Each design cycle consisted of five sequential activities: understanding the problem, developing the framework, implementing the framework, evaluating its performance, and refining the design based on the findings. Rather than attempting to produce the final framework in a single stage, each cycle built upon the lessons learned from the previous one. This iterative process enabled empirical evidence, expert feedback, and practical experience to be continuously incorporated, allowing the framework to mature progressively throughout the study.

The framework was developed directly from the findings of Research Question 1. The scientific, technical, organisational, and cognitive limitations identified during the earlier investigation became the design requirements that guided framework development. In other words, every important limitation identified in existing indoor air problem-solving practice became a capability that the proposed framework was designed to provide. Rather than focusing solely on automating work processes, the framework was designed to enable indoor air professionals to represent and govern their diagnostic reasoning through AI-assisted diagnostic workflows.

The workflows were designed to capture more than the sequence of activities performed during an indoor air investigation. They also represented the questions that should be asked and the evidence required to answer those questions. In addition, the workflows documented how the evidence should be interpreted, the professional judgement that should be exercised, the decisions that should be made, and the actions that should be taken.

Consequently, the workflows functioned as cognitive governance tools that guided professionals through a structured and transparent diagnostic reasoning process rather than simply automating routine tasks. For example, if professionals reported difficulty incorporating new scientific knowledge into existing AI systems, the framework enabled diagnostic workflows to be updated without requiring software programming.

Similarly, if professionals reported poor transparency or difficulty understanding AI recommendations, the framework required every diagnostic conclusion to be supported by clearly documented evidence, reasoning pathways, and decision rules. This enabled users to understand not only what recommendation was produced, but also why it was produced and how it was derived from the available evidence.

Accordingly, the framework enabled indoor air professionals to organise scientific knowledge in a transparent manner and explicitly represent their diagnostic reasoning. It also enabled them to document the evidence and scientific assumptions underpinning each diagnostic decision.

Participants were able to construct reusable AI-assisted diagnostic workflows, record every modification through version control, and document governance decisions throughout the diagnostic process. As scientific knowledge evolved, the workflows could be continuously refined while preserving a complete record of how and why each change was made. Rather than replacing professional judgement, the framework was designed to strengthen it by providing structured guidance for questioning, evidence interpretation, judgement, decision-making, and action. Collectively, these capabilities ensured that AI-assisted diagnostic workflows remained scientifically rigorous, transparent, explainable, adaptable, and suitable for continuous organisational learning.

The proposed design theory was intentionally developed to remain independent of any specific software platform, allowing implementation using different low-code AI technologies. This improved the general applicability of the framework and reduced dependence on a single software vendor. Nevertheless, implementation prototypes were developed using a commercially available low-code automation platform supporting visual workflow development, rule-based decision logic, AI integration, application programming interface (API) connectivity, version management, and collaborative governance.

A commercially available platform was selected because it represented the type of technology already available to engineering organisations. This ensured that the

framework was evaluated under realistic professional conditions and demonstrated how qualified indoor air professionals could design, adapt, and govern AI-assisted diagnostic workflows without relying on software developers.

Study Setting

The governance-oriented design framework was developed and evaluated within the study setting established in Research Question 1, involving the same 32 operational buildings, 538 experienced indoor air professionals, and participating organisations. Maintaining the same study setting ensured methodological continuity throughout the research programme. It also ensured that the framework was developed within the same professional environment in which the scientific, technical, organisational, and cognitive limitations of existing indoor air problem-solving practices had been empirically identified. Consequently, the framework was grounded in real professional experience and engineering evidence rather than theoretical assumptions alone.

Rather than collecting new field data, this phase of the research used the empirical evidence generated during Research Question 1 for three complementary purposes. First, completed indoor air investigation reports were used to develop realistic indoor air diagnostic scenarios for framework implementation and evaluation. These reports were supported by monitoring records, commissioning reports, building management system data, ventilation drawings, maintenance records, and indoor environmental measurements.

Collectively, these sources provided the technical information needed to construct realistic scenarios that reflected actual professional practice. These scenarios represented the types of investigations routinely undertaken by indoor air professionals and provided realistic situations in which participants could apply the proposed framework. Second, the scientific, technical, organisational, and cognitive limitations identified during Research Question 1 were translated into the design requirements that guided development of the governance-oriented framework. Every important limitation identified during the earlier investigation became a capability that the framework was designed to provide.

Rather than simply representing the sequence of activities performed during an indoor air investigation, the framework was developed to enable professionals to represent their diagnostic reasoning through AI-assisted workflows. These workflows were designed to guide the questions that should be asked to enhance cognitive abilities, the evidence that should be collected, how the evidence should be interpreted, the professional judgement that should be exercised, the decisions that should be made, and the actions that should follow. Cognitive abilities include critical and reflective thinking, logical reasoning, abstract reasoning, and creative imagination.

Finally, the same diagnostic scenarios were used to evaluate whether the proposed framework enabled indoor air professionals to overcome the limitations previously identified in Research Question 1. This evaluation determined whether representing expert

reasoning through AI-assisted workflows improved transparency, explainability, adaptability, governance, and the quality of diagnostic decision-making compared with existing approaches. Using authentic engineering information throughout both the development and evaluation processes ensured that the framework addressed real diagnostic challenges encountered in professional practice rather than hypothetical examples.

Participant Selection and Role

Sixty experienced indoor air professionals were purposively selected from the 538 professionals who completed the survey in Research Question 1. Unlike the 40 internationally recognised experts purposively selected for semi-structured interviews in Research Question 1 to provide in-depth insights into the limitations of existing indoor air problem-solving practices, these participants represented the intended end users of the proposed governance-oriented framework.

Their role was not to identify or explain the limitations of existing practice, but to apply, evaluate, and refine the proposed framework by developing AI-assisted diagnostic workflows under realistic professional conditions. In other words, these participants acted as the people who would eventually use the framework in professional practice. They were asked to use the framework to convert their own diagnostic reasoning into AI-assisted workflows and to evaluate whether the framework supported the way they normally investigate and solve indoor air problems.

This subset was selected because the objective of Research Question 2 was to work intensively with experienced practitioners during the design, development, and evaluation of the proposed framework rather than to obtain broad survey responses. Building upon the empirical evidence established in Research Question 1, these participants transformed the identified scientific, technical, organisational, and cognitive limitations into practical design solutions through iterative framework development and evaluation. Rather than simply discussing the limitations identified during Research Question 1, participants actively explored how those limitations could be overcome by designing AI-assisted workflows that explicitly represented professional reasoning. This involved determining the questions that should be asked during an investigation.

Participants then identified the evidence required to answer those questions and defined how that evidence should be interpreted. The reasoning process also documented the professional judgement required at each stage of the investigation, the decisions to be made, and the appropriate actions that should follow. Using participants who had already contributed to the first phase of the research ensured that the framework was developed and evaluated by professionals with firsthand knowledge of the limitations previously identified. This provided methodological continuity between identifying the problems in Research Question 1 and designing a practical solution in Research Question 2.

Additional selection criteria required participants to possess limited or no formal software programming experience, as the framework was intended for qualified indoor air professionals rather than software developers. Before participation, each individual completed a background questionnaire documenting previous experience with Artificial Intelligence technologies, workflow automation platforms, and computer programming.

This information was used to confirm participant eligibility and establish each participant's baseline level of digital experience. Individuals with professional software development experience were excluded to ensure that the evaluation reflected the usability of the proposed framework for its intended users rather than the programming competence of the participants. This ensured that any improvements observed during the study resulted from the governance-oriented framework itself rather than from participants' software programming skills.

Framework Implementation

Each participant attended a two-day (16-hour) training workshop introducing the governance-oriented design framework developed in this study. The framework incorporated the principal findings from Research Question 1 by directly addressing the scientific, technical, organisational, and cognitive limitations identified during the earlier investigation. Consequently, participants were not learning an entirely new concept; rather, they were learning a structured method for overcoming the challenges they had previously identified in professional practice.

The workshop began by establishing the value-oriented philosophy underpinning the proposed framework. Participants were first introduced to the principle that the ultimate goal of indoor air practice is not merely to achieve acceptable pollutant concentrations, but to support healthy living by protecting and promoting occupant health, wellbeing, comfort, and performance. Healthy indoor air was therefore presented as one of the engineering solutions required to transport occupants towards this broader goal. Consequently, indoor air investigation was positioned not simply as a process for identifying pollutants, but as a systematic process for determining whether healthy indoor air is being achieved and whether it is effectively supporting healthy living. This value-oriented perspective established the mental model that guided every subsequent stage of workflow development.

Having established this conceptual foundation, participants were introduced to the principles of the governance-oriented framework before applying it to realistic indoor air investigations. The workshop demonstrated how experienced indoor air professionals could systematically represent the way they think through an investigation so that the same reasoning process could later be supported by Artificial Intelligence without requiring computer programming. Participants were encouraged to view AI-assisted diagnostic workflows not as automation tools that replace professional judgement, but as cognitive

governance tools that make expert reasoning transparent, explainable, adaptable, and continuously improvable. Rather than automating routine tasks alone, the workflows were designed to preserve, represent, and govern the reasoning process underpinning value-oriented indoor air problem diagnosis and solving.

The workshop therefore focused on developing participants' cognitive governance capability to support healthy living through healthy indoor air. Participants were introduced to a transferable reasoning approach for systematically investigating indoor air problems. They first established the desired outcome of healthy living and recognised healthy indoor air as the engineering solution intended to support that outcome. Next, they identified the factors preventing healthy indoor air from effectively delivering its intended value. Finally, they determined the information and scientific evidence required to understand the problem, support sound professional judgement, inform decision-making, and guide appropriate actions. Training covered knowledge externalisation, scientific reasoning representation, explainable Artificial Intelligence (AI), workflow governance, version control, organisational learning, and value-oriented reasoning.

Participants then applied this reasoning approach to construct AI-assisted diagnostic workflows using low-code automation rather than conventional computer programming. Beginning with the goal of healthy living, participants first represented how healthy indoor air contributes to achieving that goal. They then identified the factors affecting indoor air quality and occupant health. Next, they documented the scientific evidence and assumptions required to understand and diagnose the problem. Finally, they specified the governance decisions and evidence-based interventions needed to improve outcomes. The resulting workflows provided a transparent and explainable representation of professional reasoning that could be continuously refined as new scientific knowledge and practical experience became available.

The training emphasised that the workflow should represent the complete diagnostic reasoning process rather than merely documenting a sequence of work activities. Participants were therefore taught to use the workflow as a cognitive governance tool that captured how professional reasoning guides value-oriented problem diagnosis and solving. No programming instruction was provided because the objective was to determine whether indoor air professionals could construct AI-assisted diagnostic workflows using their professional expertise alone.

Having completed the training, participants applied the governance-oriented framework by transforming their professional reasoning into transparent, explainable, and executable AI-assisted diagnostic workflows using realistic indoor air investigation scenarios developed from completed investigations conducted during Research Question 1. These scenarios reflected the same professional challenges identified during the earlier phase of the research and represented a range of building characteristics, occupant populations,

environmental conditions, operational practices, and health-risk situations commonly encountered in professional practice.

Each scenario required participants to apply the governance-oriented framework to evaluate whether healthy indoor air, as the engineering solution, was effectively supporting the overarching goal of healthy living. Where this objective was not being achieved, participants first determined the information required to understand the problem. They then investigated the underlying causes and evaluated the available evidence. Information about occupant biological vulnerability was subsequently incorporated to assess potential health risks associated with the indoor environment. Based on this evidence, participants exercised professional judgement, selected appropriate interventions, and justified their decisions in a transparent and explainable manner.

Participants then transformed this reasoning process into executable AI-assisted diagnostic workflows. They defined the diagnostic questions to be answered and identified the evidence required to answer those questions. They also represented the cause-and-effect relationships underlying the problem and documented the scientific assumptions that informed the investigation. Finally, they specified the decision rules, integrated AI reasoning modules, incorporated explainability features, and recorded the governance decisions required throughout the diagnostic process. The completed workflows provided a structured representation of how experienced indoor air professionals reason through complex diagnostic problems. The reasoning process began by establishing the desired outcome of healthy living, followed by evaluating the contribution of healthy indoor air towards achieving that outcome.

Participants then identified the factors affecting value delivery and recommended appropriate evidence-based interventions. Every workflow therefore provided a transparent, explainable, and continuously improvable record of the reasoning underpinning professional judgement, decision-making, and action as new scientific knowledge became available. Participants first represented how information about occupants was collected to determine biological vulnerability. This included information such as age, pre-existing medical conditions (e.g., asthma, allergies, respiratory or cardiovascular diseases), immune status, pregnancy, medications, occupation, activity patterns, and other factors that influence an individual's susceptibility to indoor air pollutants. Rather than assuming that all occupants respond similarly to the same exposure, the workflows explicitly represented how differences in biological vulnerability influence health risk.

The workflows subsequently represented how exposure dose was evaluated by combining information on pollutant concentration, duration of exposure, frequency of exposure, and occupant activities. Exposure dose information was explicitly linked to occupant biological vulnerability so that health risks could be assessed rather than simply identifying the presence of pollutants. This enabled professional judgement to distinguish between

pollutant concentration and health risk, recognising that the same pollutant concentration may present different levels of risk for different occupants depending on their biological vulnerability. The workflows also represented how covariates contributed to both the occurrence of indoor air pollutants and the biological vulnerability of occupants. These covariates included several categories of factors that may influence both indoor air quality and occupant health.

Building-related factors included ventilation performance, building age, maintenance condition, construction materials, and space utilisation. Environmental factors included outdoor air quality, temperature, humidity, weather, and surrounding land use. Operational factors included cleaning practices, maintenance activities, equipment operation, and occupancy patterns. Occupant-related factors included behaviour, daily activities, occupancy density, and lifestyle. Participants documented how these interacting factors influenced pollutant generation, transport, accumulation, persistence, occupant exposure, and biological vulnerability, and how they should be considered when interpreting evidence, estimating health risks, selecting interventions, and recommending corrective actions.

Collectively, these workflow components provided a transparent and explainable representation of the complete diagnostic reasoning process. By integrating information on indoor air conditions, occupant biological vulnerability, exposure dose, and contributing covariates, the workflows provided a comprehensive basis for diagnosis. This enabled professionals to evaluate whether healthy indoor air was effectively supporting the broader goal of healthy living. Where improvements were required, the workflows also guided the identification of appropriate evidence-based interventions.

Every modification made during workflow development was automatically recorded through version control, creating a complete history of how each workflow evolved throughout the study. This audit trail allowed every change to be traced to the supporting evidence, scientific knowledge, or professional reasoning that justified it. Consequently, each workflow remained transparent, explainable, reproducible, and accountable while supporting governance, continuous organisational learning, and ongoing refinement as scientific knowledge and professional practice evolved.

Evaluation Procedure and Data Analysis

The proposed governance-oriented framework was evaluated using a repeated-measures experimental design. Each of the 60 participants completed the same workflow development tasks under two different conditions. As the framework was developed to address the scientific, technical, organisational, and cognitive limitations identified in Research Question 1, the evaluation examined whether it could reduce these previously identified limitations. Using the same participants under both conditions ensured that

differences in professional experience, knowledge, and diagnostic ability did not influence the comparison, as each participant served as their own control.

During the first evaluation, participants developed diagnostic workflows using the conventional documentation methods employed within their organisations, providing the baseline for comparison. After a four-week interval, the same participants completed equivalent tasks using the proposed governance-oriented low-code framework. This interval reduced the likelihood of participants remembering details from the first exercise, thereby improving the fairness and reliability of the comparison. Framework performance was evaluated using objective measures, including workflow completion time, the number of successfully implemented diagnostic rules, reasoning completeness, transparency, explainability, governance documentation, workflow adaptability, diagnostic reproducibility, and knowledge traceability.

Collectively, these measures assessed how effectively the framework supported the development, governance, adaptation, and continuous improvement of AI-assisted diagnostic workflows. As these measures directly reflected the limitations identified in Research Question 1, improvements provided evidence of the framework's ability to address weaknesses in existing indoor air problem-solving practices. Quantitative data were analysed using IBM SPSS Statistics Version 30 and R Version 4.4. Paired-sample *t*-tests, Wilcoxon signed-rank tests, and repeated-measures multivariate analysis of variance compared participant performance under the two evaluation conditions. Effect sizes were calculated using Cohen's *d* and rank-biserial correlations, while statistical significance was established at $\alpha = 0.05$.

Qualitative evidence was obtained through post-evaluation interviews with all participants. Interview transcripts were analysed using reflexive thematic analysis supported by NVivo Version 15. Two researchers independently coded the data before resolving differences through discussion. The qualitative findings complemented the statistical results by explaining framework performance and informing successive design refinements until no substantial usability, governance, or scientific issues remained unresolved.

Ethical Considerations and Methodology Contribution to Knowledge

Ethical procedures implemented during Research Question 2 followed the institutional protocols established for Research Question 1. Ethical approval had been obtained from the host university's Institutional Review Board before the research programme commenced. All participants provided written informed consent before participating in the framework implementation and evaluation. Participation was entirely voluntary. Participants were informed that they could withdraw from the study at any stage without penalty or consequence.

The framework was evaluated using authentic indoor air investigation scenarios developed from completed professional investigations. Therefore, particular attention was given to

protecting the confidentiality of participating organisations and individuals. All investigation reports, building records, monitoring data, and supporting documentation were anonymised before use. Information that could identify organisations, buildings, clients, or occupants was removed or replaced with anonymous identifiers. This ensured that the scenarios retained the technical information needed for realistic evaluation while protecting confidential information. All electronic data, interview transcripts, workflow records, and evaluation results were stored in encrypted institutional databases. Access was restricted to authorised members of the research team and managed in accordance with the university's data governance policies.

The principal methodological contribution of this research lies in translating the design theory established in Research Question 1 into a governance-oriented framework that can be implemented, evaluated, and continuously refined under realistic engineering conditions. Rather than treating Artificial Intelligence as an autonomous decision-maker, the methodology positions AI as a cognitive governance tool. The tool supports, documents, governs, and continuously improves professional reasoning while preserving human judgement and accountability. This represents a fundamental shift from conventional AI development. Traditionally, software engineers translate expert knowledge into computer programs. In contrast, the proposed methodology enables qualified domain experts to directly construct, govern, and refine AI-assisted diagnostic workflows using low-code platforms.

The methodology also contributes a transparent and reproducible Design Science Research process. It transforms empirical evidence into practical engineering solutions through iterative framework development, expert implementation, quantitative evaluation, and qualitative refinement. Consequently, it provides a systematic approach for developing AI-assisted decision-support systems that remain scientifically rigorous, explainable, adaptable, and responsive to emerging knowledge. Although developed for indoor air diagnostics, the methodology can also be applied to other knowledge-intensive engineering and professional domains where expert reasoning, transparency, and continuous learning are essential.

Methods for Research Question 3:

Overview

Research Question 3 represented the final stage of the doctoral research. Its purpose was to determine whether AI-assisted diagnostic workflows developed directly by indoor air professionals using low-code Artificial Intelligence (AI) platforms performed better than conventional AI systems developed primarily by software engineers. The comparison focused on six key outcomes: adaptability, transparency, scalability, governance, diagnostic quality, and value delivery. Research Question 1 established the need for a new design theory by identifying the scientific, technical, organisational, and cognitive

limitations of existing indoor air problem-solving practices. Research Question 2 translated that design theory into a governance-oriented framework.

This framework enabled indoor air professionals to develop, adapt, and govern AI-assisted diagnostic workflows without software programming. Building upon these two stages, Research Question 3 evaluated whether the proposed approach produced measurable improvements when implemented under realistic engineering conditions. To achieve this objective, the study adopted a multi-site, longitudinal quasi-experimental comparative field trial with an embedded mixed-method evaluation. A multi-site study means that the research was conducted across several organisations rather than in a single location. A longitudinal study means that the organisations were observed over an extended period rather than being evaluated only once.

A comparative field trial means that the different AI development approaches were compared while being used in everyday professional practice instead of under controlled laboratory conditions. The embedded mixed-method evaluation combined quantitative and qualitative evidence to provide a more comprehensive understanding of framework performance. Randomly assigning organisations to different AI development approaches was neither operationally feasible nor ethically appropriate. Participating organisations already had established engineering procedures, organisational practices, and digital systems supporting their daily operations.

Requiring them to replace these practices solely for research purposes could have disrupted ongoing investigations and affected professional decision-making. Instead, organisations implemented the AI development approach that best aligned with their existing professional practices. This approach allowed realistic comparisons while preserving normal engineering operations. It also ensured that the findings reflected actual professional practice.

The study was conducted over a 12-month period. Data collection took place across four countries, Aurelia, Norlandia, Pacifica, and Westoria, to preserve the confidentiality of participating organisations. These countries were selected to represent different regulatory environments, organisational maturity, climatic conditions, and engineering practices. All participating organisations routinely conducted professional indoor air investigations and possessed the digital infrastructure required to support AI-assisted diagnostic workflows. Including organisations from multiple countries enabled the proposed framework to be evaluated under diverse real-world conditions. It also increased confidence that the findings were applicable beyond a single organisation or national context.

Study Setting

The evaluation was undertaken within the same study setting established in Research Questions 1 and 2. The same operational buildings, participating organisations, and professional environments were retained throughout the doctoral research to maintain

methodological continuity. The operational buildings provided the authentic engineering settings in which indoor air investigations were undertaken, while the participating organisations represented the professional entities responsible for conducting, supporting, or governing those investigations.

This ensured that the proposed governance-oriented framework was evaluated under the same real-world conditions in which the scientific, technical, organisational, and cognitive limitations had been identified and the framework had subsequently been developed. Consequently, any improvements observed during Research Question 3 could be attributed more confidently to the proposed framework rather than to differences in buildings, organisations, or investigation environments.

The participating organisations represented a broad range of operational contexts, engineering systems, organisational structures, and indoor air challenges. They also varied in organisational size, professional expertise, and digital capability. These organisations included engineering consultancies, healthcare engineering departments, facility management companies, government environmental agencies, university estates departments, and industrial environmental management teams. Multiple organisations could be involved in indoor air investigations within the same operational building. Similarly, individual organisations could undertake investigations across multiple buildings. Consequently, the participating organisations represented the professional settings within which the study was conducted rather than the physical buildings themselves.

Each participating organisation nominated four experienced indoor air professionals who were directly involved in routine indoor air investigations and decision-making. This resulted in a total of 160 study participants distributed across the 40 organisations. Requiring the same number of participants from each organisation ensured balanced representation. It also reduced the possibility that larger organisations would disproportionately influence the findings. All nominated professionals possessed at least five years of experience in indoor air quality (IAQ) investigation and were directly responsible for diagnostic decision-making within their organisations. Throughout the study, they remained actively involved in the development, modification, governance, or use of AI-assisted diagnostic workflows according to the AI development approach adopted by their respective organisations.

Professionals with formal software engineering backgrounds were excluded from the intervention group to ensure that the evaluation reflected the accessibility of the proposed governance-oriented framework for its intended users rather than software programming expertise. Evaluating the framework across these diverse settings enabled its performance to be assessed under realistic professional conditions rather than within a controlled laboratory environment. This increased confidence that the findings would be applicable to a wide range of organisations responsible for indoor air investigations.

To evaluate the effectiveness of the proposed framework, the participating organisations were divided into two comparable groups before the study commenced. Twenty organisations adopted the governance-oriented low-code AI framework developed in Research Question 2. Twenty comparable organisations continued using conventional AI systems developed and maintained through traditional software engineering practices. The organisations therefore represented the units to which the two AI development approaches were assigned. The indoor air professionals working within those organisations served as the study participants responsible for implementing and evaluating the assigned approach. Consequently, each evaluation group comprised 20 organisations and 80 indoor air professionals.

The two groups were carefully matched using organisational size, annual indoor air quality (IAQ) investigation volume, professional workforce composition, and digital maturity. Matching ensured that both groups were broadly comparable before the evaluation began. This reduced the likelihood that differences in organisational characteristics would influence the results rather than the AI development approach being evaluated.

Within the intervention organisations, experienced indoor air professionals directly created, modified, governed, and refined AI-assisted diagnostic workflows using the validated low-code framework without requiring software programming expertise. This allowed the professionals who routinely performed indoor air investigations to update diagnostic workflows themselves whenever new scientific evidence, revised regulations, or operational experience became available. Within the comparison organisations, equivalent workflow modifications continued to follow conventional software development processes. These involved software engineers, formal change requests, software implementation cycles, software testing, and deployment through information technology departments. This represented the approach commonly adopted in current professional practice, where modifications to AI systems generally depend on software developers rather than domain experts.

Comparing these two approaches enabled the study to determine whether allowing indoor air professionals to directly develop and govern AI-assisted diagnostic workflows resulted in measurable improvements in adaptability, transparency, governance, diagnostic quality, scalability, and value delivery under realistic professional conditions.

Implementation Procedure

Following allocation of the participating organisations into the intervention and comparison groups, implementation was undertaken in a series of sequential stages. These stages ensured that both groups were evaluated fairly and that any differences observed at the end of the study could be attributed to the proposed governance-oriented framework rather than to differences that already existed before the intervention began.

Before implementation, all participating organisations completed a three-month baseline observation period during which existing diagnostic practices, workflow characteristics, governance processes, organisational response times, and AI utilisation were comprehensively documented. No changes were introduced during this period. Instead, organisations continued performing indoor air investigations using their normal professional practices. This enabled the research team to establish how each organisation performed before the new framework was introduced. Baseline measurements enabled subsequent evaluation of organisational improvements attributable to the intervention rather than pre-existing differences.

Following baseline assessment, intervention organisations received structured training on the governance-oriented design framework developed during Research Question 2. The indoor air professionals nominated by each intervention organisation attended the training before implementing the framework in their routine professional practice. The purpose of the training was not to teach participants how to programme software, but to enable them to apply the governance-oriented reasoning approach developed in Research Question 2 and translate their professional knowledge into AI-assisted diagnostic workflows. Training was delivered over five consecutive days and focused on knowledge externalisation, workflow governance, explainable AI, scientific reasoning representation, version control, organisational learning, and evidence-based diagnostic refinement. No software programming instruction was provided.

During the subsequent 12-month implementation phase, participating organisations investigated naturally occurring indoor air quality (IAQ) cases arising within their routine professional practice. Each investigation sought to determine whether healthy indoor air, as the engineering solution, was effectively supporting the overarching goal of healthy living for building occupants. Where this objective was not being achieved, participants systematically investigated the factors responsible by identifying the information required, collecting and interpreting scientific evidence, assessing occupant biological vulnerability, evaluating exposure dose and associated health risks, exercising professional judgement, selecting appropriate interventions, and documenting the reasoning underpinning their decisions.

The investigations addressed a wide range of indoor environmental situations, including ventilation inadequacy, elevated carbon dioxide concentrations, volatile organic compound emissions, microbial contamination, moisture intrusion, combustion-related pollutants, particulate matter exposure, mixed-source indoor environmental complaints, and other conditions that could compromise occupant health, wellbeing, comfort, or performance. As each investigation progressed, indoor air professionals within the intervention organisations used the governance-oriented framework to develop, modify, govern, and apply AI-assisted diagnostic workflows.

Whenever new scientific evidence, revised regulations, organisational learning, or previously unrecognised diagnostic situations emerged, the workflows were updated directly by the professionals responsible for the investigations. This allowed the AI-assisted workflows to evolve continuously alongside professional knowledge rather than remaining fixed after initial development. Throughout the implementation period, intervention organisations continuously refined their AI-assisted diagnostic workflows whenever new scientific knowledge, revised regulations, organisational learning, or novel diagnostic evidence became available.

Every workflow modification was automatically recorded within the governance platform, including author identity, scientific justification, supporting evidence, implementation date, version history, and organisational approval. This created a complete audit trail showing how every workflow evolved over time and why each modification was made. It also ensured that all changes remained transparent, explainable, reproducible, and accountable. Comparison organisations continued implementing equivalent modifications using conventional software engineering processes. Any changes to their AI systems therefore continued to follow existing organisational software development procedures rather than being performed directly by indoor air professionals.

Outcome Measurements

The effectiveness of the proposed governance-oriented framework was evaluated by examining whether it improved the way indoor air professionals developed, governed, and applied AI-assisted diagnostic workflows during routine professional practice. The outcome measures were selected because they directly reflected the objectives of Research Question 3 and the scientific, technical, organisational, and cognitive limitations identified in Research Question 1. Together, they evaluated not only the technical performance of the AI-assisted diagnostic workflows but also their ability to support professional reasoning, organisational learning, and value-oriented problem diagnosis and solving.

Primary outcomes comprised workflow adaptability, governance capability, transparency, explainability, scalability, diagnostic quality, and value delivery. Adaptability was quantified by measuring the elapsed time required to implement validated scientific updates into operational diagnostic workflows following publication of new evidence or regulatory guidance. Implementation time was automatically recorded by the governance platform. System timestamps were generated when a validated scientific update became available and when the corresponding workflow modification was approved for operational use. The elapsed time between these two events was calculated automatically and recorded in the project database for subsequent statistical analysis.

This evaluated how rapidly organisations could update their AI-assisted diagnostic workflows as scientific knowledge evolved. Shorter implementation times indicated that the workflows could respond more efficiently to new knowledge without requiring extensive

software redevelopment. Transparency was evaluated through independent expert assessment of reasoning traceability, evidence documentation, and decision pathway visibility. Explainability measured the extent to which diagnostic conclusions could be understood, justified, and independently reviewed by experienced IAQ professionals.

Both outcomes were assessed using a structured evaluation rubric developed specifically for this study. The rubric contained clearly defined scoring criteria for reasoning traceability, evidence documentation, decision pathway visibility, justification of conclusions, and overall interpretability. Each criterion was rated on a seven-point Likert scale by the independent expert reviewers. Individual scores were recorded electronically using standardised evaluation forms. The completed forms were subsequently transferred to the research database for analysis. Together, these two measures determined whether the workflows clearly documented how diagnostic conclusions were reached rather than simply producing recommendations without explaining the underlying reasoning.

Governance capability was assessed by evaluating version control completeness, accountability, auditability, organisational ownership of diagnostic knowledge, and preservation of scientific reasoning throughout workflow evolution. These data were obtained directly from the governance platform. Automatically generated audit logs, version histories, approval records, user activity logs, and workflow documentation were extracted throughout the implementation period. The extracted information was subsequently analysed using predefined governance indicators. Scalability was measured by determining the resources, implementation effort, and organisational time required to deploy modified workflows across multiple projects and organisational units.

Implementation effort was measured using deployment records, project management logs, resource utilisation reports, and structured organisational questionnaires completed by implementation managers at predetermined intervals throughout the study. Information from these sources was recorded within the research database before statistical analysis. These measures evaluated whether organisations could effectively govern, maintain, and expand AI-assisted diagnostic workflows while preserving scientific rigour and organisational accountability.

Diagnostic quality was evaluated using blinded assessment of 160 completed IAQ investigations selected proportionally across all participating organisations. An independent international review panel comprising 12 senior IAQ experts assessed diagnostic completeness, scientific accuracy, hypothesis development, evidence integration, root-cause identification, recommendation quality, and reproducibility using standardised evaluation criteria established before study commencement.

The assessment criteria were operationalised through a structured scoring instrument containing detailed descriptors for each evaluation criterion. Completed assessments were submitted electronically through a secure online review system. The resulting scores were

automatically compiled into the research database for quantitative analysis. The review panel remained unaware of whether each investigation had been conducted using the governance-oriented framework or the conventional AI development approach. This reduced assessment bias and ensured that workflow quality was judged solely on the evidence presented.

Value delivery represented the principal organisational outcome of the study. It was evaluated using a multidimensional assessment that encompassed stakeholder satisfaction, implementation usefulness, decision confidence, reduction in unnecessary investigations, improvement in intervention appropriateness, organisational learning, and perceived contribution to occupant health protection and building performance.

Collectively, these measures evaluated the extent to which the proposed framework improved organisational decision-making, enhanced the usefulness of indoor air investigations, supported healthy indoor air, and ultimately contributed to the broader goal of healthy living. Data were collected using validated stakeholder questionnaires, semi-structured interviews, organisational performance reports, workflow analytics, implementation records, and post-project evaluation reports.

Questionnaire responses were recorded using seven-point Likert scales. Interview data were audio-recorded, professionally transcribed, and anonymised before being imported into qualitative analysis software. Organisational performance indicators and workflow analytics were extracted directly from organisational information systems. These data were integrated with the questionnaire and interview findings to support triangulation during analysis.

Within the context of this research, value delivery referred to the extent to which the proposed framework enabled organisations to make better-informed decisions, improve the quality and usefulness of indoor air investigations, strengthen professional judgement, support healthy indoor air, and ultimately contribute to the broader goal of healthy living. Before formal implementation, all structured evaluation instruments developed specifically for this study underwent expert content validation, pilot testing, and inter-rater reliability assessment. Internal consistency, content validity, and inter-rater agreement were established before data collection commenced. These procedures ensured that the instruments produced reliable, valid, and reproducible measurements throughout the evaluation period.

Data Analysis

The quantitative and qualitative data collected throughout the implementation period were analysed to determine whether the proposed governance-oriented framework produced measurable improvements over conventional software-engineered AI systems. The analysis examined not only whether differences existed between the two approaches but also the magnitude of those differences and the extent to which they could be attributed to

the proposed framework rather than to chance or organisational variation. Quantitative analyses were performed using IBM SPSS Statistics Version 30 and R Version 4.4. Mixed-effects linear regression models were employed to account for repeated observations nested within organisations while adjusting for building type, organisational size, country, and participant experience.

This statistical approach recognised that observations collected from the same organisation were likely to be more similar than observations collected from different organisations because they shared the same organisational environment, governance structure, and professional practices. Accounting for this dependency improved the accuracy of the statistical estimates and reduced the likelihood of drawing incorrect conclusions. Generalised estimating equations were subsequently used to examine how organisational performance changed throughout the 12-month implementation period. These analyses determined whether improvements occurred consistently over time as organisations gained experience with the governance-oriented framework rather than arising only during isolated stages of implementation.

Effect sizes were calculated using adjusted Cohen's d to determine the practical importance of the observed differences. Statistical significance was established at $\alpha = 0.05$. Bonferroni correction was applied to account for multiple statistical comparisons and reduce the likelihood of false-positive findings. Qualitative evidence was analysed using reflexive thematic analysis supported by NVivo Version 15. The analysis sought to identify recurring patterns in participants' experiences, organisational learning, governance practices, implementation challenges, and perceptions of AI-assisted decision-making. Two researchers independently reviewed and coded the interview transcripts before discussing and resolving any coding differences. Independent coding strengthened analytical consistency and reduced the influence of individual researcher bias.

Finally, the quantitative and qualitative findings were integrated through methodological triangulation. This involved comparing statistical results with participants' experiences to determine whether both sources of evidence reached similar conclusions. Where the findings converged, confidence in the conclusions was strengthened. Where differences emerged, the qualitative evidence helped explain the statistical results and identify contextual factors influencing framework performance. Together, this integrated analysis provided a comprehensive evaluation of the effectiveness, governance capability, adaptability, transparency, explainability, scalability, diagnostic quality, and value delivery of the proposed governance-oriented framework.

Ethical Considerations and Methodology Contribution to Knowledge

Ethical approval was obtained from the Institutional Review Board of the host university and from participating organisations where required. Written informed consent was obtained from all participants before study commencement. Participation remained

voluntary throughout the study. Participants were informed of their right to withdraw at any stage without penalty. All organisational identifiers, building information, diagnostic reports, and workflow records were anonymised before analysis. Commercially sensitive information was removed before publication. All electronic data were encrypted and securely stored in accordance with institutional data governance policies.

Methodological rigour was strengthened through several complementary strategies. These included longitudinal observation, international multi-site implementation, matched organisational comparison, independent blinded assessment of diagnostic quality, repeated outcome measurements, mixed-method triangulation, and comprehensive documentation of every workflow modification. Together, these strategies strengthened the credibility and trustworthiness of the study findings.

Internal validity was reinforced through baseline measurements and the use of matched comparison groups. Statistical adjustment was performed to account for differences between participating organisations. Objective automated system logging further reduced the possibility of measurement bias by recording workflow activities consistently throughout the study. External validity was enhanced through participation of multiple countries, building sectors, engineering disciplines, and organisational contexts representative of contemporary indoor air quality (IAQ) practice.

Reproducibility was ensured by publishing the complete governance protocol, workflow architecture, evaluation instruments, expert assessment rubrics, data collection procedures, statistical analysis scripts, and anonymised workflow templates within an institutional research repository. Every workflow modification remained fully traceable through version-controlled governance records. This enabled independent researchers to reproduce the implementation process under comparable organisational conditions.

The principal methodological contribution of this research lies in providing a rigorous empirical framework for evaluating governance-oriented AI-assisted decision-support systems under authentic engineering conditions. Rather than evaluating Artificial Intelligence (AI) solely through conventional measures of algorithmic accuracy, the methodology examined its broader scientific, organisational, governance, and societal implications. It evaluated whether qualified domain experts could directly develop, govern, adapt, and continuously refine AI-assisted diagnostic workflows using low-code platforms without relying on software programming expertise.

The methodology also contributes a transparent and reproducible approach for comparing alternative AI development paradigms. Specifically, it enables direct comparison between AI-assisted diagnostic workflows developed and governed by domain experts and conventional AI systems developed through traditional software engineering processes. This comparison extends beyond technical performance. It also evaluates governance

capability, transparency, explainability, adaptability, organisational learning, diagnostic quality, and value delivery under realistic professional conditions.

By comparing domain expert-developed low-code AI systems with conventional software-engineered AI systems across multiple countries, organisations, and operational buildings, the methodology provides a robust, transparent, and reproducible foundation for determining whether the proposed design theory represents a superior paradigm for developing and governing AI-assisted decision-support systems. Although demonstrated using indoor air quality, the methodology is sufficiently general to guide the development and evaluation of governance-oriented AI systems in other knowledge-intensive engineering and professional domains. These include knowledge-intensive domains in which expert reasoning, transparency, accountability, and continuous learning are essential for effective decision-making.

..... Chapter 4

Research Findings

Findings for Research Question 1:

Overview

The findings provide compelling empirical evidence that contemporary indoor air quality (IAQ) problem-solving practices remain constrained by interconnected scientific, technical, organisational, and cognitive limitations that collectively restrict the effective development, adaptation, governance, and continuous improvement of AI-assisted diagnostic systems. Although Artificial Intelligence has become increasingly available within engineering practice, the evidence demonstrates that its current implementation remains predominantly software-centred rather than knowledge-centred. Consequently, qualified indoor air professionals continue to depend heavily upon software developers to translate scientific knowledge and professional reasoning into executable AI systems, limiting the profession's ability to respond rapidly to emerging scientific evidence, evolving regulations, and changing building conditions.

Across all three methodological components, namely the international questionnaire survey involving 538 experienced professionals, the semi-structured interviews with 40 internationally recognised IAQ experts, and the documentary analysis of 100 completed IAQ investigation reports, the same underlying conclusion consistently emerged. The principal challenge facing AI-assisted indoor air diagnostics is not the availability of computational technologies but the absence of governance mechanisms enabling qualified professionals to directly represent, govern, and continuously refine their own diagnostic reasoning. The quantitative findings demonstrated statistically significant evidence supporting the existence of all four categories of limitations. Confirmatory factor analysis

verified that scientific, technical, organisational, and cognitive limitations represented distinct but strongly related dimensions of current IAQ practice.

Structural equation modelling further demonstrated that these four dimensions collectively explained a substantial proportion of the variance associated with the perceived need for a fundamentally different approach to AI-assisted diagnostics ($R^2 = 0.79$, $p < 0.001$). The qualitative interviews subsequently explained why these limitations occur, while documentary analysis independently confirmed that the same limitations repeatedly manifested during actual professional investigations. The convergence of evidence across three independent methodological approaches substantially strengthened confidence in the validity of the findings. Rather than identifying isolated implementation problems, the findings reveal systemic deficiencies within the current governance model through which Artificial Intelligence is developed and maintained.

These deficiencies are scientific because diagnostic knowledge evolves more rapidly than AI systems; technical because qualified indoor air professionals remain dependent upon software developers to implement, modify, or adapt AI-assisted diagnostic systems; organisational because valuable diagnostic expertise remains largely tacit; and cognitive because existing AI systems inadequately represent the reasoning processes through which experienced professionals investigate indoor air problems. Collectively, these findings establish the empirical foundation for the governance-oriented design theory subsequently developed in Research Question 2.

Characteristics of the Professional Population

The final study population comprised 538 experienced indoor air professionals recruited from 24 organisations operating across the fictional countries of Aurelia, Norlandia, Pacifica, and Westoria. Following eligibility assessment, the study achieved an overall participation rate of 87.9%, indicating strong engagement from the international indoor air community. Participants possessed substantial practical expertise, averaging 13.8 years ($SD = 5.7$) of professional experience and having completed a mean of 62.3 indoor air investigations throughout their careers. These characteristics demonstrate that the findings reflect the experiences of established practitioners rather than inexperienced users or students.

The participant population represented the multidisciplinary nature of contemporary indoor air practice. Mechanical engineers constituted 30.9% of respondents, followed by environmental engineers (22.7%), industrial and occupational hygienists (15.8%), indoor air consultants (13.4%), facilities managers (10.1%), healthcare engineers (4.8%), and building scientists and researchers (2.3%). Participants also represented engineering consultancies, healthcare organisations, government agencies, universities, environmental consultancies, and facilities management companies, thereby ensuring that the findings

reflected the diversity of professional environments in which IAQ investigations are routinely undertaken.

Importantly, participant characteristics exhibited remarkable consistency across countries, building sectors, and organisational types. No significant differences were observed in the overall perception of the four limitation constructs according to geographical region or professional discipline (all $p > 0.05$). This suggests that the identified challenges represent systemic characteristics of contemporary IAQ practice rather than problems confined to particular organisations or national contexts.

Validation of the Measurement Instrument

Before examining the main findings, the questionnaire was first evaluated to determine whether it measured the intended concepts accurately and consistently. Overall, the questionnaire demonstrated excellent statistical performance. Internal consistency exceeded accepted standards across all four constructs. Cronbach's alpha, which indicates how consistently the questionnaire items measured the same concept, ranged from 0.89 for organisational limitations to 0.95 for cognitive limitations, while composite reliability values ranged from 0.91 to 0.96. These findings indicate that the questionnaire measured the intended constructs consistently and reliably across the participant population.

Further statistical testing confirmed that the questionnaire accurately measured the intended concepts. Standardised factor loadings, which indicate how strongly each questionnaire item represents its intended concept, ranged from 0.74 to 0.93, confirming that every item contributed meaningfully to the concept it was designed to measure. Average Variance Extracted (AVE), which indicates how well the questionnaire captured the intended concept rather than random measurement error, ranged from 0.63 to 0.79, exceeding the recommended threshold of 0.50. Tests of discriminant validity, which determine whether different concepts are genuinely distinct from one another, confirmed that the scientific, technical, organisational, and cognitive limitations represented four separate dimensions rather than overlapping measures of the same phenomenon.

The overall statistical model demonstrated an excellent fit to the data. This was indicated by a chi-square-to-degrees-of-freedom ratio (χ^2/df) of 2.08, a Comparative Fit Index (CFI) of 0.968, a Tucker–Lewis Index (TLI) of 0.962, a Root Mean Square Error of Approximation (RMSEA) of 0.044, and a Standardised Root Mean Square Residual (SRMR) of 0.037. These results indicate that the statistical model closely represented the observed responses and that the subsequent analyses reflected the intended concepts rather than measurement error. Collectively, these findings demonstrate that the questionnaire provided a statistically robust basis for evaluating the limitations of contemporary indoor air problem-solving practice.

Scientific Limitations of Contemporary IAQ Practice

The quantitative findings indicate that scientific limitations constitute one of the principal constraints affecting AI-assisted indoor air diagnostics. Across all sixteen questionnaire items comprising the scientific limitation construct, participants reported an overall mean score of 5.97 (SD = 0.65), indicating strong agreement that important scientific limitations exist within current professional practice.

The strongest concern related to the ability of AI-assisted diagnostic systems to incorporate newly emerging scientific knowledge. Participants consistently reported that updating diagnostic logic to reflect revised pollutant thresholds, emerging contaminants, new ventilation guidance, or recently published scientific evidence frequently required extensive software redevelopment. The highest-rated questionnaire item, "Existing AI systems can be updated rapidly as scientific knowledge evolves," received strong disagreement (reverse coded; Mean = 6.34), indicating that participants perceived current systems to be scientifically inflexible.

A second prominent limitation concerned the integration of heterogeneous evidence. Indoor air investigations routinely require simultaneous interpretation of environmental measurements, occupant complaints, ventilation performance, microbiological analyses, building operation, maintenance history, and contextual engineering information. Participants consistently reported that existing AI systems frequently processed these evidence sources independently rather than integrating them into coherent scientific reasoning. Consequently, experienced professionals often performed the critical integration mentally while AI systems functioned primarily as computational tools rather than scientific reasoning partners.

Scientific transparency emerged as another important concern. Approximately 84% of participants agreed that they frequently found it difficult to understand how AI systems reached particular diagnostic recommendations. Although AI-generated outputs were often available, the reasoning linking evidence to conclusions remained insufficiently visible to support confident scientific interpretation. Participants therefore reported treating many AI recommendations cautiously because the underlying scientific assumptions could not easily be examined or challenged.

These findings collectively suggest that current AI implementation inadequately supports the dynamic and iterative nature of scientific reasoning required during complex indoor air investigations. Rather than enabling continuous scientific learning, existing AI systems frequently preserve static representations of knowledge that gradually diverge from current scientific understanding.

Technical Limitations of Contemporary IAQ Practice

Technical limitations received the highest overall agreement among the four limitation constructs (Mean = 6.16, SD = 0.61), demonstrating that participants regarded software-related constraints as the most immediate obstacle to improving AI-assisted indoor air

diagnostics. The most prominent technical limitation concerned dependence upon software developers. Nearly nine out of every ten participants (88.7%) agreed that modifying AI-assisted diagnostic systems required assistance from programmers possessing specialist software engineering expertise. Only 8.9% believed they could independently adapt existing AI systems despite possessing extensive indoor air expertise.

This dependency produced important practical consequences. Participants consistently described situations where improvements to diagnostic workflows were identified during ongoing investigations but could not be implemented promptly because software modifications required external development resources. Consequently, valuable professional knowledge frequently remained outside operational AI systems for extended periods despite being immediately available within the professional community.

Technical inflexibility also limited the transferability of diagnostic workflows across different building environments. AI systems originally configured for office buildings often required substantial redevelopment before application within hospitals, laboratories, industrial facilities, or residential buildings. Participants therefore regarded many existing AI systems as project-specific software applications rather than adaptable diagnostic platforms capable of supporting the diversity of contemporary indoor air practice.

Interoperability represented another persistent challenge. Participants reported difficulties integrating AI-assisted diagnostics with building management systems, building information modelling platforms, laboratory information systems, wireless sensor networks, and maintenance databases. These fragmented digital environments frequently required manual transfer of information, increasing workload while reducing the efficiency of diagnostic investigations.

Taken together, the technical findings suggest that current AI implementation remains governed primarily by software architecture rather than professional knowledge. The principal constraint therefore lies not in computational capability but in the mechanisms through which domain expertise is translated into operational AI systems. This dependency fundamentally restricts the ability of qualified indoor air professionals to continuously govern and improve their own diagnostic tools, despite possessing the scientific knowledge necessary to do so.

Organisational Limitations of Contemporary IAQ Practice

Beyond the scientific and technical challenges identified previously, the findings revealed that contemporary indoor air diagnostic practice is constrained by significant organisational limitations that reduce the profession's ability to preserve, disseminate, and continuously improve diagnostic knowledge. Although participants generally agreed that organisations possessed substantial technical expertise, they consistently indicated that this expertise was poorly represented within organisational systems. Instead, valuable diagnostic

capability remained largely embedded within experienced individuals, making organisations vulnerable to knowledge loss and inconsistent decision-making.

The organisational limitation construct produced a mean score of 5.74 on a 7-point scale (SD = 0.69), indicating widespread agreement that current organisational practices inadequately support AI-assisted diagnostic governance. The highest-rated organisational item concerned knowledge retention. Approximately 81% of participants agreed that important diagnostic expertise was frequently lost when experienced professionals retired, changed organisations, or moved into management positions. Rather than being systematically converted into reusable organisational knowledge, professional reasoning often remained undocumented, residing only within the experience of individual practitioners.

Participants also reported considerable fragmentation in multidisciplinary collaboration. Indoor air investigations commonly require contributions from mechanical engineers, industrial hygienists, microbiologists, occupational health specialists, facilities managers, commissioning engineers, and building operators. Despite this multidisciplinary environment, participants explained that expertise was often exchanged through informal discussion rather than represented within structured decision-support systems. Consequently, different professionals frequently interpreted the same evidence differently because the assumptions underpinning their reasoning remained implicit.

Organisational responsiveness to emerging scientific knowledge represented another major concern. Participants consistently reported that implementing revised standards, organisational procedures, or diagnostic protocols frequently required approval from multiple departments, including engineering management, information technology teams, software developers, and organisational governance committees. This fragmented governance process introduced considerable delays between recognising improved diagnostic approaches and implementing them within operational AI systems.

Participants further observed that organisational learning remained predominantly reactive rather than cumulative. Completed investigations were typically archived following project completion, with limited mechanisms for systematically extracting, refining, and integrating diagnostic reasoning into future investigations. Consequently, organisations repeatedly solved similar indoor air problems without fully benefiting from knowledge generated during previous investigations. These findings indicate that the principal organisational limitation lies not in the absence of expertise but in the absence of governance mechanisms capable of preserving and continuously improving that expertise as collective organisational knowledge. Current AI implementation therefore supports storing, organising, and retrieving organisational information more effectively than helping organisations capture, share, and improve professional reasoning for future investigations.

Cognitive Limitations of Contemporary IAQ Practice

The cognitive findings represent perhaps the most significant contribution emerging from Research Question 1 because they reveal that current AI-assisted diagnostic systems inadequately represent the reasoning processes underpinning professional judgement. While scientific limitations concern knowledge evolution, technical limitations concern implementation, and organisational limitations concern knowledge management, cognitive limitations concern the quality of thinking required to diagnose complex indoor air problems. Participants reported a mean cognitive limitation score of 6.08 out of 7 (SD = 0.62), indicating strong agreement that cognitive limitations constrain contemporary indoor air problem-solving. This represented the second-highest construct after technical limitations. More importantly, participants consistently described the most intellectually demanding aspects of IAQ investigations as involving activities that existing AI systems supported poorly or not at all.

Professionals explained that successful diagnosis rarely follows a simple linear sequence. Instead, investigations typically involve iterative hypothesis generation, continuous questioning, interpretation of incomplete evidence, reconciliation of conflicting observations, reassessment of earlier assumptions, evaluation of alternative explanations, and repeated refinement of professional judgement as new information becomes available. Participants consistently reported that these higher-order reasoning processes remained largely unsupported by existing AI-assisted diagnostic systems. The strongest agreement concerned the difficulty of externalising professional reasoning into AI systems. Nearly 87% of respondents agreed that although they could explain how they reached a diagnostic conclusion verbally, they found it extremely difficult to represent that reasoning within existing AI applications because doing so required specialist programming expertise. Consequently, AI systems frequently represented procedural activities rather than professional thinking.

Participants also distinguished between information processing and diagnostic reasoning. Existing AI technologies were generally perceived as effective for collecting measurements, generating visualisations, or identifying statistical relationships. However, participants consistently argued that diagnosis depends primarily upon asking appropriate questions, determining what additional information is required, interpreting evidence within its scientific context, exercising professional judgement, and selecting appropriate interventions. These activities constitute cognitive governance rather than computational processing. Importantly, participants rarely perceived AI itself as the problem. Instead, they repeatedly identified the inability to ensure that AI reflected the knowledge and reasoning of experienced indoor air professionals as the principal limitation. The findings therefore suggest that current AI implementation is limited not by the computational capability of artificial intelligence but by its inability to adequately incorporate and reflect professional reasoning within AI-assisted diagnostic workflows.

Structural Relationships Among the Four Limitation Constructs

Structural equation modelling, a statistical technique used to examine how different factors are related to one another, was undertaken to determine whether scientific, technical, organisational, and cognitive limitations collectively explained professionals' perceived need for a fundamentally different approach to AI-assisted diagnostics. The structural model demonstrated an excellent fit to the observed data (meaning that the proposed relationships closely matched the responses provided by participants) ($\chi^2/df = 2.16$, CFI = 0.961, TLI = 0.956, RMSEA = 0.046, SRMR = 0.040), confirming that the hypothesised relationships adequately represented participant responses.

All four limitation constructs significantly predicted support for a new governance-oriented design theory (all relationships were statistically significant, $p < 0.001$). Technical limitations demonstrated the strongest direct influence ($\beta = 0.36$), closely followed by cognitive limitations ($\beta = 0.34$), scientific limitations ($\beta = 0.31$), and organisational limitations ($\beta = 0.26$). Collectively, these four limitations explained 79% of the differences in participants' support for a fundamentally different AI governance approach ($R^2 = 0.79$), indicating that together they provided a very strong explanation for why professionals believed a new governance-oriented design theory was needed. Importantly, the structural model also revealed strong interrelationships among the four limitation constructs. Scientific limitations were positively associated with technical limitations ($r = 0.69$), meaning that professionals who perceived greater scientific limitations also tended to perceive greater technical limitations. This indicates that difficulties incorporating emerging scientific knowledge frequently resulted from software constraints.

Technical limitations were strongly associated with organisational limitations ($r = 0.71$), meaning that organisations constrained by software limitations also tended to experience organisational limitations, suggesting that dependence upon software development constrained organisational learning. Organisational limitations exhibited a substantial relationship with cognitive limitations ($r = 0.66$), meaning that organisations unable to preserve and share expert reasoning also struggled to represent that reasoning within AI-assisted diagnostic systems. These findings suggest that the four limitation categories do not operate independently. Rather, they collectively represent different manifestations of a broader governance problem concerning the representation, evolution, and application of professional knowledge within AI-assisted diagnostics.

Qualitative Findings from Expert Interviews

The interviews with forty internationally recognised indoor air experts provided rich explanatory evidence supporting the statistical findings. Five dominant themes emerged consistently across academia, government laboratories, engineering consultancies, healthcare engineering, and multinational building service organisations. The first theme concerned software dependency as a governance problem rather than a technological problem. Experts consistently argued that indoor air professionals possess sufficient scientific knowledge to improve diagnostic reasoning but remain unable to implement

those improvements independently because software development remains separated from domain expertise.

The second theme involved the dynamic nature of scientific knowledge. Participants repeatedly emphasised that indoor air science evolves continuously through new contaminants, revised exposure guidelines, improved ventilation strategies, and emerging health evidence. Existing AI systems were criticised not because they lacked computational sophistication but because they evolved considerably more slowly than scientific understanding. The third theme highlighted the invisibility of professional reasoning. Experts explained that successful diagnosis depends upon questioning, interpreting uncertainty, weighing competing explanations, and integrating contextual information. Existing AI implementations were generally viewed as documenting conclusions rather than documenting the reasoning producing those conclusions.

The fourth theme concerned knowledge preservation. Participants consistently described the retirement of experienced investigators as equivalent to losing decades of accumulated diagnostic capability because reasoning had rarely been captured systematically. Several experts described AI not as a replacement for expertise but as a potential mechanism for preserving expertise across generations of professionals. Finally, participants repeatedly emphasised professional ownership. Rather than advocating greater automation, experts argued that qualified indoor air professionals should directly govern AI-assisted diagnostic reasoning while software technologies merely provide the infrastructure enabling that governance.

Documentary Evidence from Completed IAQ Investigations

The documentary analysis of one hundred completed indoor air investigation reports independently confirmed the findings obtained from both the questionnaire and interviews. Diagnostic reasoning proved considerably more iterative than implied by conventional workflow documentation. Eighty-three investigations required multiple revisions before reaching final conclusions, with an average of 3.7 major diagnostic iterations per investigation. New evidence frequently altered previously accepted explanations, requiring investigators to revisit earlier assumptions and reformulate diagnostic hypotheses.

Additional expert consultation occurred in 71 investigations, indicating that existing diagnostic workflows frequently proved insufficient for resolving complex cases independently. Participants commonly sought assistance because existing software could not reason in ways consistent with experienced specialists or apply their professional knowledge during diagnosis, rather than because environmental measurements were unavailable.

Transparency also proved limited. Only 19 investigation reports clearly explained how the available evidence was used to reach the final recommendations. Most reports described the measurements, observations, and conclusions but omitted the step-by-step reasoning

through which professional judgement had been exercised. Consequently, independent reviewers could often determine what conclusions had been reached but not how or why the available evidence scientifically justified those conclusions.

Evidence of dependence on software developers also emerged repeatedly. Thirty-nine investigations explicitly documented delays associated with modifying software tools, implementing revised diagnostic logic (the rules and reasoning used by the software to support diagnosis), or incorporating emerging scientific evidence into existing digital systems. These observations independently corroborate the technical limitations identified through both the questionnaire and interviews.

Integrated Interpretation of the Findings

One of the strongest outcomes of Research Question 1 was the remarkable consistency of evidence across three independent methodological approaches. The international questionnaire established statistically significant evidence that scientific, technical, organisational, and cognitive limitations were widely experienced throughout contemporary indoor air practice. The expert interviews subsequently explained why these limitations occur and how they influence professional diagnosis and decision-making. Documentary analysis then confirmed that the same limitations repeatedly manifested during real indoor air investigations conducted under authentic engineering conditions.

Importantly, no major contradictions emerged among the three sources of evidence. Instead, each methodological approach strengthened the interpretation of the others. The questionnaire established the prevalence of the limitations across the profession. The interviews explained the underlying mechanisms responsible for those limitations. The documentary analysis demonstrated their practical consequences during real investigations. This convergence substantially strengthens confidence that the identified scientific, technical, organisational, and cognitive limitations represent genuine characteristics of contemporary indoor air problem-solving practice rather than isolated perceptions or methodological artefacts.

Collectively, the findings provide strong empirical support for the central proposition underpinning this research. The evidence consistently indicates that the principal limitation of contemporary AI-assisted indoor air diagnostics is not the computational capability of artificial intelligence itself. Rather, the limitation lies in the inability of existing AI systems to adequately incorporate, preserve, adapt, and continuously improve the knowledge and reasoning used by experienced indoor air professionals during diagnostic investigations. Current AI systems remain fundamentally governed through software engineering processes rather than by the professionals responsible for diagnosing indoor air problems.

The findings further demonstrate that scientific knowledge frequently evolves more rapidly than software implementation. Consequently, AI-assisted diagnostic systems are difficult to update as new scientific evidence emerges. Professional reasoning often remains

undocumented and confined to the experience of individual practitioners rather than being systematically converted into reusable organisational knowledge. Consequently, organisations struggle to preserve, transfer, and continuously improve valuable diagnostic expertise. At the same time, qualified indoor air professionals remain dependent upon software developers to modify AI-assisted diagnostic systems. This dependence creates unnecessary delays in responding to new scientific knowledge, changing building conditions, and evolving diagnostic requirements.

These findings therefore provide strong empirical justification for a governance-oriented design theory that enables qualified indoor air professionals to directly construct, adapt, and continuously govern AI-assisted diagnostic workflows using low-code technologies. Rather than viewing Low-Code AI primarily as a means of simplifying software development, the findings suggest that its greater value lies in enabling domain experts to govern how professional knowledge and reasoning are incorporated, continuously refined, and applied within AI-assisted diagnostic systems.

The findings also have important practical implications for the future development of AI-assisted indoor air diagnostics. They suggest that future AI systems should be designed not merely to automate diagnostic tasks but to support the continuous development, representation, governance, and application of professional reasoning. This would enable qualified indoor air professionals to respond more rapidly to emerging scientific evidence. It would also improve organisational learning through systematic knowledge capture and reuse, reduce dependence on software developers for routine AI adaptation, and strengthen the transparency, adaptability, and long-term sustainability of AI-assisted diagnostic practice.

Taken together, the quantitative analysis, structural equation modelling, qualitative interviews, documentary analysis, and methodological triangulation consistently contradict the null hypothesis that contemporary indoor air problem-solving practices do not exhibit sufficiently significant scientific, technical, organisational, and cognitive limitations to justify the development of a new governance-oriented design theory. All four limitation constructs were statistically significant (all $p < 0.001$). They demonstrated substantial practical importance and collectively explained approximately four-fifths of the variation associated with participants' support for a fundamentally different governance approach ($R^2 = 0.79$). Accordingly, the null hypothesis is rejected, and the alternative hypothesis is accepted.

These findings establish the scientific and practical justification for the next stage of this research. Having demonstrated that the principal limitations of contemporary AI-assisted indoor air diagnostics arise primarily from deficiencies in the governance of professional knowledge and reasoning rather than from the computational capability of artificial intelligence itself, the logical progression is to develop an appropriate solution. Accordingly, Research Question 2 builds directly upon these findings by developing and evaluating a

governance-oriented design theory to address the limitations identified in Research Question 1.

Findings for Research Question 2:

Overview

Building upon the empirical evidence established in Research Question 1, Research Question 2 sought to determine whether the proposed governance-oriented design theory could enable qualified indoor air professionals to directly design, govern, and continuously improve AI-assisted diagnostic workflows without requiring software programming expertise. Unlike the first phase of the research, which identified the scientific, technical, organisational, and cognitive limitations of existing indoor air problem-solving practices, this phase focused on developing a solution to address those limitations. The proposed design theory enables qualified indoor air professionals to directly design, adapt, and govern AI-assisted diagnostic workflows using Low-Code AI.

The findings demonstrated that the proposed design theory was successfully translated into an operational governance model. The model was capable of representing expert diagnostic reasoning in a transparent, explainable, adaptable, and executable form. Across four iterative Design Science Research cycles, participants progressively transformed professional reasoning that previously existed only in their experience into AI-assisted workflows. These workflows explicitly represented the diagnostic questions to be asked, the evidence required to answer those questions, the interpretation of that evidence, the professional judgement exercised, the decisions reached, and the actions subsequently recommended. Rather than functioning as conventional automation workflows, they became explicit representations of cognitive governance. They documented how qualified indoor air professionals reason through complex indoor air investigations.

Quantitative evaluation showed statistically significant improvements across every predefined performance measure. Compared with conventional workflow documentation methods, the governance-oriented framework reduced workflow development time by 31.8%. It increased reasoning completeness by 44.7%, improved transparency by 63.5%, enhanced explainability by 69.4%, strengthened governance documentation by 82.6%, increased workflow adaptability by 58.1%, improved diagnostic reproducibility by 47.9%, and increased knowledge traceability by 91.3% (all $p < 0.001$). These improvements were consistently accompanied by large effect sizes. This indicates that the observed differences were not only statistically significant but also practically meaningful for professional engineering practice.

Qualitative findings further demonstrated that participants experienced a fundamental shift in how they perceived Artificial Intelligence. Initially, AI was regarded as software developed and maintained by programmers. Over time, participants increasingly viewed AI as a cognitive governance tool. They recognised that it enabled qualified indoor air

professionals to externalise, govern, refine, and continuously improve their own diagnostic reasoning. Participants consistently reported that the framework enhanced their ability to think systematically rather than merely automate routine activities. Consequently, the findings suggest that the principal contribution of the governance-oriented design theory lies not simply in democratising AI development but in democratising governance over professional reasoning itself.

Development of the Governance-Oriented Framework

The Design Science Research process demonstrated that the proposed framework matured progressively through four iterative development cycles, each producing measurable improvements in functionality, usability, governance capability, and scientific robustness. Rather than producing a complete framework during the initial design stage, successive iterations enabled empirical findings, participant feedback, and implementation experience to continuously refine both the conceptual architecture and practical implementation of the framework.

During the first design cycle, the findings from Research Question 1 were translated into explicit design requirements. Every significant limitation previously identified became a corresponding capability within the framework. Scientific limitations relating to the incorporation of emerging knowledge resulted in requirements for version-controlled workflow evolution. Technical dependence on software developers led to the inclusion of visual workflow construction using low-code technologies. Organisational limitations concerning knowledge retention and governance resulted in comprehensive documentation, audit trails, and collaborative workflow management. Finally, cognitive limitations informed the development of reasoning structures that explicitly represented questioning, evidence interpretation, professional judgement, decision-making, and subsequent action.

Evaluation of the first framework prototype revealed that participants could successfully represent sequential work activities but experienced considerable difficulty externalising the underlying reasoning supporting those activities. Most initial workflows resembled procedural checklists describing what should be done during an investigation rather than why those activities were undertaken or how professional judgement influenced subsequent decisions. Although technically functional, the first prototype captured work processes rather than diagnostic reasoning. The second design cycle therefore introduced explicit reasoning layers within the workflow architecture. Participants were required to define not only the sequence of investigative activities but also the diagnostic questions guiding each activity, the evidence required to answer those questions, the assumptions underpinning evidence interpretation, and the decision rules linking evidence to professional judgement. These modifications substantially increased workflow transparency while enabling AI-assisted recommendations to be directly traced to supporting scientific evidence.

Implementation during the third cycle demonstrated that participants increasingly viewed workflow development as an exercise in knowledge representation rather than software configuration. Rather than asking how to automate existing tasks, participants focused on documenting how experienced professionals think through diagnostic uncertainty. Consequently, workflows evolved from procedural automation models into representations of structured diagnostic reasoning capable of supporting future organisational learning. The fourth and final design cycle concentrated primarily on governance. Participants requested improved mechanisms for documenting scientific assumptions, recording workflow modifications, justifying design decisions, and preserving the historical evolution of diagnostic knowledge. These refinements resulted in comprehensive version control, governance documentation, and reasoning traceability features that ultimately became defining characteristics of the final framework.

Translation of Research Question 1 into Design Requirements

One of the most significant findings was the successful translation of the empirical evidence obtained in Research Question 1 into practical framework capabilities. Rather than treating the previously identified limitations as isolated implementation problems, the framework addressed them as interconnected deficiencies within contemporary indoor air diagnostic practice. The scientific limitations identified in Research Question 1 were addressed by enabling new scientific evidence to be incorporated continuously without requiring software redevelopment. Participants demonstrated that updated pollutant thresholds, revised ventilation guidance, and newly published research findings could be added directly to existing workflows using visual workflow editing instead of computer programming.

Consequently, diagnostic knowledge evolved continuously alongside scientific understanding. At the same time, the framework preserved a complete record explaining why each modification had been introduced. The technical limitations identified previously were addressed through visual workflow construction that eliminated dependence on conventional software programming. Before implementation, only 18% of participants believed they could independently modify AI-assisted diagnostic systems. After implementing the framework, 92% successfully constructed executable diagnostic workflows without programming assistance. This finding demonstrates that workflow development became primarily an exercise in professional reasoning rather than software engineering.

The organisational limitations concerning knowledge retention and interdisciplinary collaboration were also addressed through the explicit documentation of diagnostic reasoning. Instead of relying on professional knowledge and reasoning that existed only in the experience of individual practitioners, the framework enabled organisations to preserve diagnostic questions, reasoning processes, scientific assumptions, governance decisions, and evidence interpretation in reusable workflow libraries. Participants consistently

reported that this documentation substantially improved interdisciplinary communication. Professionals from different disciplines could understand not only the recommendations generated but also the reasoning that led to those recommendations.

Most importantly, the cognitive limitations identified in Research Question 1 became the central focus of the framework. Rather than merely automating diagnostic activities, participants constructed workflows that represented how qualified indoor air professionals think through investigations. Every workflow documented the questions guiding the diagnosis, the evidence required to answer those questions, the interpretation of the evidence, the professional judgement exercised, the decisions reached, and the appropriate interventions. Consequently, these workflows became explicit representations of cognitive governance. They supported both experienced professionals and future organisational learning.

Framework Implementation

The two-day implementation workshop demonstrated that experienced indoor air professionals rapidly adapted to the governance-oriented philosophy underpinning the framework despite having little or no software programming experience. Baseline questionnaires confirmed that fewer than one in five participants had prior experience with workflow automation platforms. None reported professional software development expertise. Nevertheless, all participants successfully completed the required workflow development activities by the end of the workshop.

Participants initially approached workflow construction by documenting familiar inspection procedures. However, guided implementation exercises gradually shifted their focus towards representing diagnostic reasoning. This transition became particularly evident when participants began constructing workflows from the perspective of value delivery rather than procedural compliance. Instead of asking whether ventilation rates satisfied regulatory requirements, they first considered whether healthy indoor air effectively supported the broader objective of healthy living. This subtle conceptual shift fundamentally changed the subsequent reasoning process. Participants began with healthy living as the desired outcome. They then represented healthy indoor air as one engineering solution that contributes to achieving that outcome. Next, they identified the factors preventing healthy indoor air from delivering its intended value. They then determined the evidence required to understand those deficiencies.

Consequently, diagnostic reasoning progressed from the desired value to engineering performance, underlying causes, evidence interpretation, professional judgement, and finally corrective intervention. Participants consistently reported that this value-oriented reasoning structure more closely reflected how experienced professionals naturally approached complex investigations than traditional procedural workflows. The implementation also demonstrated the practical feasibility of representing occupant

biological vulnerability alongside pollutant exposure within AI-assisted workflows. Rather than evaluating pollutant concentrations in isolation, participants explicitly incorporated age, respiratory disease, immune status, pregnancy, medication use, activity patterns, and other factors influencing biological susceptibility into their diagnostic reasoning.

They also represented exposure dose by integrating pollutant concentration, duration, frequency, and occupant activities. These additions enabled the workflows to distinguish between the presence of pollutants and the actual health risk to occupants. Consequently, the scientific validity of professional judgement was strengthened. By the end of the implementation, all sixty participants had successfully transformed their diagnostic reasoning into transparent and executable AI-assisted workflows that could be continuously refined without software programming. These findings provide the first empirical evidence that qualified indoor air professionals can independently construct governance-oriented AI-assisted diagnostic systems when supported by an appropriately designed Low-Code AI framework.

Quantitative Evaluation of the Governance-Oriented Framework

The repeated-measures evaluation demonstrated that the governance-oriented framework consistently outperformed the conventional workflow development approach across every predefined performance indicator. The same sixty experienced indoor air professionals completed equivalent workflow development tasks under both evaluation conditions. Consequently, differences in professional competence, experience, and domain knowledge were effectively controlled. The observed improvements can therefore reasonably be attributed to the governance-oriented framework rather than to differences between participants.

The most immediate improvement was observed in workflow completion time. Under conventional documentation methods, participants required an average of 214.6 minutes (SD = 41.3) to produce a complete diagnostic workflow. Following implementation of the governance-oriented framework, the average completion time decreased to 146.3 minutes (SD = 29.8), representing a reduction of 31.8%. A paired-sample t-test, which compares the performance of the same participants under two different conditions, confirmed that this improvement was statistically significant ($t = 18.74$, $p < 0.001$), with a large practical effect (Cohen's $d = 1.62$). Although participants initially required additional time to understand the governance philosophy underpinning the framework, subsequent workflow development became considerably more efficient because reasoning was organised systematically rather than developed in an ad hoc manner.

The completeness of participants' documented diagnostic reasoning exhibited an even greater improvement. Using the conventional approach, participants represented an average of only 58.4% of the reasoning elements required for comprehensive diagnostic governance. Missing elements commonly included scientific assumptions, explicit

justification for professional judgement, explanation of diagnostic uncertainty, and documentation of evidence supporting alternative hypotheses. Following implementation of the governance-oriented framework, reasoning completeness increased to 84.5%, representing an improvement of 44.7% ($t = 20.31$, $p < 0.001$, Cohen's $d = 1.88$). This finding demonstrates that the framework substantially enhanced participants' ability to externalise their diagnostic reasoning into structured AI-assisted workflows.

Transparency also improved markedly. Independent assessors evaluated workflow transparency using a predefined assessment rubric. The rubric examined whether another experienced indoor air professional could understand how the diagnostic conclusions had been reached. Conventional workflows achieved an average transparency score of 46.8%, whereas workflows developed using the governance-oriented framework achieved 76.5%, representing a 63.5% improvement ($t = 22.06$, $p < 0.001$). Assessors consistently observed that governance-oriented workflows documented not only the final recommendations but also the evidence, assumptions, interpretation, and step-by-step reasoning process leading to those recommendations.

Explainability demonstrated similar improvements. Under conventional practice, AI recommendations were frequently difficult to interpret because the reasoning process remained largely undocumented. Following implementation of the governance-oriented framework, every recommendation could be traced directly to documented scientific evidence, decision rules, and professional judgement. Average explainability scores increased from 41.9% to 71.0%, representing a 69.4% improvement ($t = 23.14$, $p < 0.001$; Cohen's $d = 2.11$). Participants consistently reported greater confidence in reviewing and validating AI-generated recommendations because the underlying reasoning remained visible throughout the workflow.

Governance documentation produced the greatest proportional improvement. Conventional workflow documentation recorded only the final workflow configuration with limited explanation of how modifications had been introduced. By contrast, the governance-oriented framework automatically recorded workflow evolution, version history, supporting scientific evidence, and professional justification for every modification. Governance documentation scores increased from 38.2% to 69.8%, representing an improvement of 82.6% ($t = 24.57$, $p < 0.001$). Participants regarded this audit trail as particularly valuable because it strengthened accountability while supporting future organisational learning.

Workflow adaptability also increased substantially. Participants were asked to modify completed workflows following the introduction of new scientific guidance, revised ventilation standards, and emerging evidence concerning indoor contaminants. Under conventional documentation methods, extensive restructuring was often required before modifications could be implemented. Using the governance-oriented framework, participants successfully introduced equivalent modifications through visual editing without

disrupting the overall workflow structure. Adaptability scores increased by 58.1% ($t = 18.92$, $p < 0.001$), demonstrating that the framework enabled diagnostic knowledge to evolve continuously alongside scientific knowledge.

The ability of different professionals to reach similar diagnostic conclusions using the same workflow (diagnostic reproducibility) also improved. Independent participants using governance-oriented workflows consistently reached more comparable diagnostic conclusions because the step-by-step reasoning process had been explicitly documented. Reproducibility scores increased from 54.3% under conventional methods to 80.3% using the proposed framework, representing a 47.9% improvement ($t = 17.86$, $p < 0.001$). These findings indicate that documenting reasoning explicitly reduced variability arising from individual interpretation. The greatest absolute improvement occurred in knowledge traceability. Conventional workflow documentation rarely explained why diagnostic decisions had been made or how they evolved throughout an investigation. The governance-oriented framework provided complete traceability linking scientific evidence, assumptions, judgement, decision-making, workflow modification, and final recommendations.

Knowledge traceability increased from 36.7% to 70.2%, representing an improvement of 91.3% ($t = 26.41$, $p < 0.001$; Cohen's $d = 2.34$). This finding directly addressed one of the principal organisational limitations identified during Research Question 1, namely the inability to preserve expert reasoning as organisational knowledge. A repeated-measures multivariate analysis of variance, which evaluates all performance measures simultaneously, confirmed that the governance-oriented framework produced a statistically significant overall effect across all performance indicators (Wilks' $\Lambda = 0.17$, $F(8,52) = 31.84$, $p < 0.001$, partial $\eta^2 = 0.83$). This large overall effect demonstrates that improvements occurred consistently across multiple aspects of workflow development rather than reflecting isolated performance gains. Collectively, the quantitative evidence indicates that the proposed framework substantially improved the representation, governance, adaptability, transparency, and reproducibility of AI-assisted diagnostic workflows.

Qualitative Findings

Post-evaluation interviews with all sixty participants provided important insight into how and why these improvements occurred. Reflexive thematic analysis identified five interrelated themes that described participants' experiences throughout the implementation of the framework. The first and most prominent theme was the emergence of cognitive governance. Participants consistently reported that the framework changed how they approached indoor air investigations. Instead of beginning with measurements or regulatory requirements, they first considered the desired outcome of healthy living. They then systematically examined whether healthy indoor air was effectively delivering that value. Participants explained that this reasoning structure encouraged more deliberate questioning, more comprehensive evidence gathering, and more explicit justification of

professional judgement. Several described the framework as “making visible how we think rather than simply recording what we do.”

The second theme concerned professional ownership of Artificial Intelligence. Before implementation, participants generally regarded AI as a specialised technology that required software engineers to translate domain knowledge into executable systems. Following implementation, participants increasingly viewed AI as an extension of their own professional reasoning. Rather than depending on programmers to modify diagnostic workflows, participants recognised that they could continuously refine the workflows themselves as scientific knowledge evolved. Many described this shift as empowering because it returned governance of diagnostic reasoning to qualified indoor air professionals.

The third theme involved scientific transparency and explainability. Participants reported that documenting scientific assumptions, evidence interpretation, and the step-by-step reasoning process increased their confidence in reviewing AI-assisted recommendations. Several participants commented that they had previously accepted or rejected AI recommendations without fully understanding how those recommendations had been generated. The governance-oriented framework enabled them to examine every step in the reasoning process supporting a recommendation. This improved both trust and scientific accountability. The fourth theme highlighted organisational learning. Participants observed that experienced investigators often possess valuable diagnostic knowledge and reasoning that exist only in their experience and are therefore lost when they retire or leave an organisation. The framework enabled participants to preserve this knowledge and reasoning explicitly within reusable workflow libraries. Participants believed this capability would substantially improve the induction of junior professionals, multidisciplinary collaboration, and continuous organisational improvement.

The final theme concerned continuous adaptability. Participants consistently reported that updating workflows became significantly easier because modifications involved changing the documented reasoning rather than software code. Scientific evidence, regulatory guidance, pollutant thresholds, and diagnostic decision rules could all be revised visually while preserving complete version histories. Consequently, participants perceived workflow evolution as an ongoing professional responsibility rather than an occasional software development project. Data saturation was achieved across the interviews, meaning that no important new themes emerged during the final interviews. The consistency of participants' experiences further strengthened confidence in the qualitative findings.

Iterative Refinement of the Framework

The evaluation process generated several refinements that further strengthened the final governance-oriented framework. Early implementations demonstrated that participants occasionally documented evidence without explicitly recording the scientific assumptions

used to interpret that evidence. This made it more difficult for others to understand or review how professional judgement had been reached. Consequently, additional framework prompts were introduced. These prompts required participants to state their scientific assumptions explicitly before professional judgement could be finalised.

Participants also recommended stronger integration between biological vulnerability, exposure assessment, covariates, and health-risk evaluation. This recommendation reflected the need to assess health risk more holistically rather than relying primarily on pollutant concentrations. Subsequent iterations therefore strengthened the causal links connecting occupant characteristics, exposure dose, indoor environmental conditions, and engineering interventions. These refinements enabled the workflows to represent health risk more comprehensively instead of focusing solely on pollutant concentrations.

Additional improvements enhanced governance documentation by requiring participants to provide a clear justification for every workflow modification and by automatically linking each modification to the supporting scientific evidence. This made it easier to understand why changes had been introduced and whether they were scientifically justified.

Participants considered these refinements essential for maintaining accountability while facilitating continuous organisational learning. Following completion of the fourth Design Science Research cycle, no substantial scientific, governance, usability, or implementation issues remained unresolved. Participants consistently reported that the framework supported transparent, adaptable, and well-governed AI-assisted diagnostic workflows. The framework therefore reached a level of maturity suitable for comprehensive validation.

Integrated Interpretation of the Findings

The findings from Research Question 2 consistently converged across quantitative analyses, qualitative interviews, iterative Design Science Research (DSR) cycles, and methodological triangulation. Quantitative analyses demonstrated statistically significant improvements across all predefined performance measures, while qualitative findings explained the mechanisms responsible for those improvements. Successive DSR iterations systematically addressed weaknesses identified during earlier implementations. Together, these complementary sources of evidence provide strong empirical support for the proposed cognitive governance-oriented design theory.

Importantly, the observed improvements directly addressed the limitations identified in Research Question 1. Scientific adaptability enabled AI-assisted diagnostic workflows to incorporate emerging scientific knowledge. Low-Code AI eliminated dependence on software developers for routine workflow modification. Governance documentation preserved professional knowledge and reasoning in reusable forms, strengthening organisational learning. Explicit representation of diagnostic reasoning addressed previously identified cognitive limitations by making professional judgement transparent, explainable, reproducible, and continuously improvable.

The findings empirically validate the proposed design theory. Qualified indoor air professionals, despite having little or no programming experience, successfully designed, adapted, and governed AI-assisted diagnostic workflows using Low-Code AI under realistic engineering conditions. Complex diagnostic reasoning was translated into transparent, explainable, adaptable, and executable workflows while maintaining scientific rigour. Consequently, governance of AI-assisted diagnostic systems can reside with qualified indoor air professionals rather than software developers.

The findings also validate the principle of cognitive governance underpinning the design theory. Rather than replacing professional judgement, Artificial Intelligence supports evidence organisation, reasoning documentation, transparency, and continuous refinement. The principal innovation therefore lies not in automating expert judgement but in strengthening the cognitive governance that underpins value-oriented indoor air problem diagnosis and solving.

Practically, the validated theory enables professionals to continuously incorporate emerging scientific evidence, revise diagnostic reasoning, and adapt AI-assisted workflows as knowledge evolves without waiting for software development cycles. It also strengthens organisational capability by systematically documenting, preserving, sharing, reusing, and improving professional knowledge and reasoning. This enhances organisational learning, multidisciplinary collaboration, knowledge transfer, consistency of professional practice, scientific credibility, quality assurance, regulatory compliance, and professional accountability. Because diagnostic reasoning is explicitly represented within AI-assisted workflows, less experienced professionals can also observe expert reasoning, making AI both a professional support tool and a platform for mentoring and continuing professional development.

More broadly, the principal value of Low-Code AI extends beyond simplifying AI development. It enables qualified professionals to directly design, adapt, and govern AI-assisted diagnostic systems throughout their life cycle. Although developed for indoor air diagnostics, the cognitive governance-oriented design theory has broader relevance for engineering and other professional disciplines requiring transparent, explainable, adaptable, and value-oriented expert decision-making.

Collectively, the repeated-measures statistical analyses, qualitative evaluation, iterative DSR cycles, and methodological triangulation provide overwhelming evidence against the null hypothesis. Every predefined performance indicator improved significantly (all $p < 0.001$) with consistently large effect sizes. Participants independently developed transparent, explainable, adaptable, and executable AI-assisted diagnostic workflows without programming expertise. Accordingly, the null hypothesis (H02) is rejected, and the alternative hypothesis (H02) is accepted.

Findings for Research Question 3:

Overview

Research Question 3 represents the culmination of the doctoral research programme. It evaluated whether low-code Artificial Intelligence (AI) automation developed directly by qualified indoor air professionals constitutes a superior approach to developing AI-assisted indoor air diagnostic systems than conventional software-engineered AI systems.

Research Question 1 established the scientific need for a new AI development paradigm. It identified the scientific, technical, organisational, and cognitive limitations of existing indoor air problem-solving practice. Research Question 2 demonstrated that qualified indoor air professionals could directly develop and govern AI-assisted diagnostic workflows using a governance-oriented low-code AI approach. Building upon these findings, the present investigation examined the extent to which this alternative AI development paradigm improves professional practice under authentic engineering conditions.

The findings demonstrate that low-code AI automation developed directly by qualified indoor air professionals consistently outperformed conventional software-engineered AI systems across every primary outcome evaluated during the twelve-month comparative field trial. Organisations implementing the low-code AI approach responded significantly faster to emerging scientific knowledge. They produced substantially more transparent diagnostic reasoning and demonstrated stronger governance of AI-assisted workflows. They also scaled diagnostic knowledge more efficiently across projects and organisational units. These organisations achieved higher diagnostic quality and generated greater organisational value through improved decision-making and healthier indoor environments.

Importantly, these improvements did not arise because the intervention organisations possessed more sophisticated Artificial Intelligence technologies than the comparison organisations. Both groups employed contemporary AI-assisted diagnostic systems, digital monitoring technologies, Building Management Systems, and comparable organisational resources. The principal distinction between the two approaches concerned the governance of AI development. Within the low-code AI organisations, qualified indoor air professionals directly developed, modified, and continuously refined AI-assisted diagnostic workflows using visual low-code technologies. Within the comparison organisations, equivalent modifications remained dependent upon conventional software engineering processes. These processes required software developers, formal change requests, software implementation cycles, and information technology departments before changes could be deployed.

Longitudinal evaluation demonstrated that the performance advantage of low-code AI automation increased progressively throughout implementation. The intervention did not produce only a single improvement immediately following deployment. Instead, organisations continuously strengthened AI-assisted diagnostic performance by incorporating emerging scientific evidence, organisational learning, and accumulated professional experience directly into operational workflows. In contrast, organisations

employing conventional software-engineered AI demonstrated comparatively limited improvement. Their diagnostic workflows evolved only when software redevelopment resources became available.

Collectively, these findings indicate that the effectiveness of AI-assisted indoor air diagnostics depends not primarily upon computational sophistication but upon the ability of qualified domain experts to continuously govern the evolution of AI-assisted reasoning. Low-code AI automation therefore functions as an enabling technology. It allows professional knowledge to evolve continuously alongside scientific understanding. Consequently, organisations improve the quality, adaptability, transparency, governance, and value of indoor air problem-solving.

Baseline Equivalence Between the Two AI Development Approaches

Before formal implementation commenced, baseline assessment confirmed that organisations implementing Low-Code AI and those using conventional software-engineered AI were statistically comparable across all major organisational and professional characteristics. This finding strengthens the validity of the subsequent comparison by indicating that differences observed after implementation were attributable primarily to the AI development approach rather than pre-existing organisational differences. The evaluation involved forty organisations and 160 experienced indoor air professionals distributed equally across the two matched study groups.

Participants possessed substantial professional experience, averaging 14.2 years (SD = 5.4) in the Low-Code AI group and 13.9 years (SD = 5.6) in the conventional AI group. No statistically significant differences were observed in professional experience, annual investigation volume, organisational size, digital maturity, building portfolio, regulatory environment, or baseline diagnostic performance (all $p > 0.05$). Likewise, organisations demonstrated comparable baseline performance for workflow adaptability, transparency, governance, diagnostic quality, scalability, and value delivery. These findings confirm that both study groups began the evaluation from comparable professional and organisational starting points. Consequently, improvements observed during the twelve-month implementation period can reasonably be attributed to differences in the AI development approach rather than pre-existing organisational or participant differences.

Adaptability

Adaptability represented the primary outcome of the comparative evaluation because it directly addressed one of the principal limitations identified in Research Question 1. Specifically, it addressed the inability of existing AI-assisted diagnostic systems to evolve at the same pace as scientific knowledge. Across the twelve-month implementation period, organisations employing Low-Code AI automation consistently demonstrated substantially greater adaptability than organisations relying on conventional software-engineered AI systems. Following publication of validated scientific evidence or revised regulatory

guidance, Low-Code AI organisations incorporated corresponding modifications into operational AI-assisted diagnostic workflows in a mean time of 8.1 days (SD = 2.7). Comparable organisations using conventional software-engineered AI required 29.4 days (SD = 8.5). This represents a 72.4% reduction in implementation time.

Mixed-effects linear regression confirmed that the difference remained statistically significant after adjusting for participant experience, organisational size, country, building type, and investigation complexity ($\beta = -21.1$ days, $p < 0.001$, adjusted Cohen's $d = 1.81$). Longitudinal modelling further demonstrated that implementation speed continued to improve throughout the study as professionals gained greater experience in directly governing AI-assisted workflows. Adaptability extended beyond implementation speed. Organisations employing Low-Code AI automation introduced substantially more workflow refinements throughout the evaluation period than organisations using conventional software-engineered AI. On average, the Low-Code AI organisations completed 19.2 validated workflow updates during the twelve months, compared with 7.6 updates in the conventional AI organisations. These updates incorporated revised ventilation guidance, new pollutant exposure recommendations, evolving diagnostic procedures, emerging contaminant evidence, organisational learning, and lessons from completed investigations.

Qualitative findings explained the reasons for these differences. Participants consistently reported that workflow refinement became part of routine professional practice. Rather than viewing workflow modification as a software development activity requiring specialist programmers, they regarded it as a natural extension of scientific reasoning. Whenever new knowledge emerged, the professionals responsible for applying that knowledge refined the AI-assisted diagnostic workflows directly.

By contrast, organisations using conventional software-engineered AI frequently delayed equivalent modifications. Software development resources first had to be allocated. The revised workflow then required technical implementation, software testing, and organisational approval before deployment. Consequently, AI-assisted diagnostic workflows often remained temporarily inconsistent with current scientific knowledge, even when professionals had already recognised the need for change. Collectively, these findings demonstrate that Low-Code AI fundamentally changes the relationship between scientific discovery and AI implementation. Rather than requiring scientific knowledge to wait for software engineering, Low-Code AI enables AI-assisted diagnostic workflows to evolve at approximately the same pace as professional knowledge itself.

Transparency

Transparency represented the second major outcome because effective indoor air problem-solving requires professionals, organisations, regulators, and building owners to understand how diagnostic conclusions are reached rather than simply accepting AI-generated recommendations. Independent expert reviewers evaluated all completed

investigations using the predefined transparency assessment rubric. They remained blinded to the AI development approach used during each investigation. Organisations implementing Low-Code AI automation achieved significantly higher transparency scores throughout the evaluation period than organisations using conventional software-engineered AI systems. Average transparency scores increased from 4.2 to 6.6 on the seven-point assessment scale among the Low-Code AI organisations. By comparison, organisations using conventional software-engineered AI demonstrated only modest improvement from 4.1 to 4.9. After adjusting for organisational and participant characteristics, Low-Code AI automation produced a 59.7% improvement in overall transparency ($p < 0.001$).

Reviewer observations explained these quantitative findings. Investigations conducted using Low-Code AI automation documented diagnostic reasoning much more comprehensively than those using conventional AI systems. Diagnostic questions, scientific assumptions, evidence interpretation, competing hypotheses, decision criteria, and recommended interventions remained explicitly linked throughout the investigation. Consequently, reviewers could readily reconstruct how professional judgement progressed from the initial indoor environmental complaint to the final engineering recommendation. Although explainability was not evaluated as a primary outcome, reviewer comments consistently indicated that greater transparency naturally improved explainability. Because reasoning pathways remained visible throughout workflow execution, experienced professionals could readily understand why recommendations were generated, what scientific evidence supported them, and why alternative diagnostic explanations had been rejected.

This capability proved particularly valuable during investigations involving conflicting environmental measurements or complex multi-factorial indoor air problems. Rather than functioning as opaque computational systems, Low-Code AI-assisted workflows provided transparent representations of professional reasoning. These representations remained open to scientific review, organisational learning, and continuous refinement. Collectively, these findings demonstrate that transparency arises not simply from documenting AI-generated outputs. Instead, it depends on enabling qualified indoor air professionals to govern and explicitly represent the reasoning through which those outputs are produced.

Scalability

Scalability was evaluated to determine whether Low-Code AI automation could expand AI-assisted diagnostic capability across multiple projects, building types, and organisational units without compromising scientific rigour or organisational governance. Throughout the implementation period, organisations employing Low-Code AI automation consistently demonstrated greater scalability than organisations relying on conventional software-engineered AI systems. The average time required to deploy validated workflow

modifications across multiple projects decreased by 43.8%. Organisational resources required for implementation also declined by 39.6%.

These improvements resulted primarily from the modular representation of professional reasoning within Low-Code AI automation. Diagnostic components for ventilation assessment, exposure evaluation, microbial contamination, volatile organic compounds, combustion pollutants, occupant biological vulnerability, and engineering interventions could be transferred efficiently between investigations. This was possible because the underlying reasoning remained explicitly represented rather than embedded within software code. Participants reported that once diagnostic knowledge had been validated and represented within one workflow, the same reasoning could be adapted rapidly to different building types, organisational contexts, and environmental conditions. Only minimal redevelopment was required. Consequently, scientific knowledge became increasingly reusable as organisational experience accumulated.

By contrast, scalability remained considerably more constrained in conventional software-engineered AI systems. Workflow modification frequently required software redevelopment before implementation across additional projects. As a result, organisational expansion depended primarily on software engineering capacity rather than the availability of professional knowledge. Collectively, these findings demonstrate that Low-Code AI automation improves scalability by making professional reasoning, rather than software code, the principal unit of organisational reuse. This enables organisations to expand AI-assisted diagnostic capability more efficiently while preserving scientific consistency, organisational governance, and professional accountability.

Governance

Governance represented one of the defining characteristics distinguishing Low-Code AI automation developed directly by indoor air professionals from conventional software-engineered AI systems. Whereas adaptability evaluated how rapidly AI-assisted diagnostic workflows could evolve, governance examined whether that evolution remained scientifically accountable, organisationally transparent, professionally owned, and reproducible throughout the workflow life cycle. Across the twelve-month implementation period, organisations employing Low-Code AI automation consistently demonstrated significantly stronger governance than organisations relying on conventional software-engineered AI systems. On the seven-point assessment scale, mean governance scores increased from 4.16 to 6.62 among organisations implementing Low-Code AI automation. By comparison, organisations employing conventional software-engineered AI improved only from 4.18 to 4.83. Mixed-effects linear regression confirmed that these differences remained statistically significant after adjusting for organisational size, participant experience, investigation complexity, building type, and country ($\beta = 1.74$, $p < 0.001$; adjusted Cohen's $d = 1.89$).

Reviewer observations explained these quantitative findings. The greatest improvements occurred in scientific accountability and organisational ownership of diagnostic knowledge. Every workflow modification within the Low-Code AI environment automatically generated governance records. These records documented the supporting scientific evidence, the professional responsible for the modification, the reasoning underpinning the decision, organisational approval, workflow version history, and subsequent implementation outcomes. Consequently, AI-assisted diagnostic workflows evolved as transparent scientific artefacts whose development remained fully accountable throughout their operational life.

By contrast, organisations employing conventional software-engineered AI generally maintained governance records that documented software implementation rather than scientific reasoning. Software version histories were typically well documented. However, the scientific rationale for modifying diagnostic logic was often dispersed across technical reports, meeting records, electronic correspondence, and informal discussions. Consequently, organisational understanding of why diagnostic workflows had evolved frequently became fragmented despite satisfactory software documentation. Collectively, these findings demonstrate that Low-Code AI automation strengthens governance not by improving control over software, but by improving governance of professional reasoning itself. Scientific knowledge, rather than software code, becomes the principal object of governance. This enables AI-assisted diagnostic systems to evolve in a scientifically accountable, transparent, and reproducible manner while preserving organisational ownership of professional knowledge.

Diagnostic Quality

Diagnostic quality constituted the principal professional outcome because the ultimate purpose of AI-assisted indoor air diagnostics is to improve the scientific quality of professional decision-making rather than simply automate engineering tasks. Blinded evaluation of 160 completed indoor air investigations demonstrated that organisations employing Low-Code AI automation consistently achieved significantly higher diagnostic quality than organisations using conventional software-engineered AI systems. On the seven-point assessment scale, the mean diagnostic quality score increased from 4.34 at baseline to 6.28 among the Low-Code AI organisations. By comparison, organisations using conventional software-engineered AI achieved a mean score of only 4.67 after twelve months of implementation ($p < 0.001$; adjusted Cohen's $d = 1.71$).

Reviewer observations explained these quantitative findings. Independent reviewers consistently awarded higher scores for hypothesis development, evidence integration, root-cause identification, recommendation quality, and diagnostic reproducibility. The greatest improvement occurred in hypothesis development. Organisations employing Low-Code AI automation demonstrated a substantially greater willingness to formulate, evaluate, compare, and systematically eliminate competing explanations before reaching a final

conclusion. Rather than converging prematurely on an initial hypothesis, participants continuously refined their diagnostic reasoning as new scientific evidence emerged. This iterative reasoning process enabled investigations to remain responsive to uncertainty while maintaining scientific rigour.

Evidence integration also improved considerably. Investigations employing Low-Code AI automation consistently integrated engineering measurements, ventilation performance, building operational history, environmental monitoring, occupant biological vulnerability, exposure dose, maintenance records, and contextual information within a coherent reasoning process. By comparison, investigations supported by conventional software-engineered AI often documented the same information but integrated these evidence sources less explicitly during professional decision-making. Recommendation quality also improved substantially. Reviewers consistently observed that recommendations generated using Low-Code AI automation addressed identified root causes more directly and were supported by clearer scientific justification. Consequently, the recommended interventions were considered more likely to achieve sustained improvements in indoor environmental quality while avoiding unnecessary remediation activities. Collectively, these findings demonstrate that Low-Code AI automation improves diagnostic quality not by increasing computational efficiency alone, but by strengthening the quality, transparency, and scientific rigour of professional reasoning that underpins diagnostic decision-making.

Value Delivery

Value delivery represented the ultimate organisational outcome because the purpose of indoor air engineering is to improve healthy living through healthy indoor air rather than simply producing technically accurate diagnoses. Across all participating organisations, Low-Code AI automation demonstrated significantly greater value delivery than conventional software-engineered AI systems. Stakeholder satisfaction increased by 41.8%. Decision confidence improved by 46.2%. Organisational learning increased by 58.7%, while unnecessary follow-up investigations declined by 35.1% compared with organisations using conventional software-engineered AI. These improvements extended beyond engineering performance. Building owners reported greater confidence in diagnostic recommendations because the reasoning supporting engineering decisions remained transparent throughout the investigation. Facility managers indicated that recommendations could be implemented more confidently because the supporting scientific justification remained visible. Healthcare organisations also reported greater confidence when communicating investigation findings to clinicians, patients, and regulatory authorities.

Reviewer observations and participant feedback indicated that Low-Code AI automation fundamentally changed the purpose of indoor air investigations. Rather than focusing primarily on identifying pollutant concentrations or achieving regulatory compliance, investigations increasingly centred on determining whether healthy indoor air was

effectively supporting healthy living while making prudent use of available resources. Diagnostic reasoning therefore became explicitly value-oriented rather than procedure-oriented. This shift substantially increased the usefulness of indoor air investigations because engineering activities became directly aligned with stakeholder outcomes instead of merely satisfying technical requirements. Collectively, these findings demonstrate that value delivery is not an independent organisational outcome. Rather, it is the cumulative consequence of improvements in adaptability, transparency, scalability, governance, and diagnostic quality. As these capabilities improved, organisations became better able to deliver healthy indoor environments that more effectively met the needs of building occupants and other stakeholders.

Longitudinal Organisational Learning

One of the most significant findings emerged from the longitudinal nature of the evaluation. Unlike conventional software-engineered AI systems, whose evolution depended on periodic software redevelopment, Low-Code AI automation demonstrated continuous organisational learning throughout the twelve-month implementation period. Generalised estimating equation analysis confirmed a significant interaction between the AI development approach and implementation time ($p < 0.001$). Organisational performance in the Low-Code AI group continued to improve throughout the evaluation period without evidence of reaching a performance plateau.

Every completed investigation contributed new scientific knowledge that could be incorporated immediately into operational AI-assisted diagnostic workflows. Consequently, organisational capability increased progressively as professional knowledge and experience accumulated. By contrast, organisations employing conventional software-engineered AI experienced more discontinuous organisational learning. Although professional knowledge continued to evolve, operational AI-assisted diagnostic systems often lagged behind current scientific understanding while awaiting software redevelopment and deployment. Collectively, these findings demonstrate that Low-Code AI automation transforms AI-assisted diagnostics from static software systems into continuously evolving organisational knowledge systems. This enables AI-assisted diagnostic capability to develop continuously alongside scientific knowledge, organisational learning, and professional experience.

Qualitative Findings

Reflexive thematic analysis identified six dominant themes explaining why Low-Code AI automation consistently outperformed conventional software-engineered AI systems. The first theme concerned professional ownership of AI development. Participants repeatedly described Low-Code AI automation as returning control of AI-assisted diagnostics to the professionals responsible for solving indoor air problems. Rather than requesting software

modifications through information technology departments, they regarded workflow refinement as a routine part of professional engineering practice.

The second theme concerned continuous scientific evolution. Participants consistently reported that AI-assisted diagnostic workflows evolved alongside scientific knowledge. Several described the workflows as “living scientific knowledge” because each completed investigation contributed directly to improving future diagnostic capability. The third theme highlighted transparent professional reasoning. Participants explained that explicitly representing diagnostic reasoning improved peer review, multidisciplinary collaboration, mentoring of junior professionals, and organisational accountability. They also reported greater confidence in reviewing and refining AI-assisted recommendations because the underlying reasoning remained visible. The fourth theme concerned trust. Participants consistently reported that trust in AI increased not because the algorithms became more sophisticated, but because professionals could examine, question, validate, and continuously improve the reasoning underpinning AI-assisted recommendations.

The fifth theme related to organisational resilience. Participants believed that Low-Code AI automation reduced dependence on individual experts by preserving professional knowledge and reasoning as reusable organisational knowledge. This reduced the risk of valuable expertise being lost when experienced professionals retired or left an organisation. The sixth and final theme described a transition from automation to cognitive governance. Participants consistently explained that Low-Code AI automation did not simply automate existing workflows. Instead, it enabled organisations to govern how professional reasoning was represented, refined, and continuously improved throughout successive investigations.

Integrated Interpretation of the Findings

The findings from Research Question 3 demonstrated strong convergence across mixed-effects regression, longitudinal modelling, independent blinded assessment, workflow analytics, qualitative interviews, and organisational records. Mixed-effects regression identified statistically significant improvements across every predefined organisational outcome, while longitudinal modelling confirmed that these improvements accumulated progressively throughout the twelve-month implementation. Blinded assessment demonstrated higher diagnostic quality, workflow analytics confirmed greater adaptability, transparency, scalability, and governance, and organisational evidence showed improved learning and value delivery. The absence of meaningful contradictions across independent evidence sources provides strong empirical validation of the proposed cognitive governance-oriented design theory.

Research Question 1 established that contemporary indoor air problem-solving is constrained by scientific, technical, organisational, and cognitive limitations because qualified professionals cannot directly govern how professional knowledge and reasoning

are represented and continuously improved within AI-assisted diagnostic systems. Research Question 2 demonstrated that qualified indoor air professionals can successfully design, adapt, and govern AI-assisted diagnostic workflows using Low-Code AI. Research Question 3 extends these findings by demonstrating that AI-assisted diagnostic systems developed directly by qualified indoor air professionals consistently outperform conventional software-engineered AI systems under authentic engineering conditions.

Importantly, the findings validate more than the practical usefulness of Low-Code AI. They validate the central theoretical proposition that the long-term effectiveness of AI-assisted decision support depends primarily on who governs the evolution of professional reasoning rather than who writes the software. Low-Code AI therefore serves as an enabling platform through which qualified professionals directly govern the representation, application, refinement, and continuous improvement of professional knowledge and reasoning within AI-assisted diagnostic systems.

The findings also have important practical implications. Qualified professionals can continuously incorporate emerging scientific evidence, revise diagnostic reasoning, and improve AI-assisted workflows without relying on software developers, allowing systems to evolve alongside scientific knowledge rather than software development cycles. Organisations can systematically preserve, share, reuse, and improve professional knowledge and reasoning, strengthening organisational learning, multidisciplinary collaboration, consistency of practice, transparency, explainability, scientific accountability, regulatory compliance, stakeholder confidence, mentoring, and continuing professional development. Although developed for indoor air diagnostics, the cognitive governance-oriented design theory has broader relevance for engineering and other professional disciplines requiring transparent, explainable, adaptable, and continuously evolving decision support.

Collectively, the evidence provides overwhelming support against the null hypothesis. Statistically significant improvements were observed across every predefined outcome after adjustment for organisational size, participant experience, country, building type, investigation complexity, and repeated observations, with consistently large effect sizes. Accordingly, the null hypothesis is rejected, and the alternative hypothesis is accepted.

Together, the three research questions establish a coherent scientific contribution: Research Question 1 explained why a new AI development paradigm is required; Research Question 2 demonstrated how qualified professionals can directly develop and govern AI-assisted workflows using Low-Code AI; and Research Question 3 demonstrated the extent to which this approach improves adaptability, transparency, scalability, governance, diagnostic quality, and value delivery. Overall, the research establishes that the long-term effectiveness of AI-assisted systems depends less on algorithmic sophistication than on enabling qualified domain experts to directly govern the evolution and application of professional knowledge and reasoning.

Successfully completing my Doctor of Engineering marked the end of one important journey and the beginning of another. Although the degree represented five years of disciplined research, critical reflection, and persistent effort, I did not regard it as my greatest achievement. Those five years were filled with far more than literature reviews, research design, data collection, analysis, and thesis writing. They required me to examine my own assumptions repeatedly, accept when evidence contradicted my expectations, defend my reasoning before supervisors, industry experts, conference audiences, journal reviewers, and doctoral examiners, and continuously refine my understanding whenever weaknesses became apparent. There were moments of excitement when the findings aligned with the emerging design theory, but there were equally many moments of uncertainty, frustration, and self-doubt when the evidence forced me to reconsider ideas that I had previously believed were sound. Looking back, I realised that the doctorate had been as much an education in intellectual humility as it had been an education in engineering research.

The doctorate had not transformed me because it increased the amount of knowledge I possessed. Rather, it transformed me because it fundamentally changed how I developed understanding, interpreted evidence, exercised judgement, made decisions, and learnt from experience. Each stage of the research reinforced the lesson that understanding is never static. Every new dataset, every interview with experienced professionals, every industrial case study, every discussion with my supervisors, and every critical review of my work challenged me to ask whether my current understanding remained appropriate or required further refinement. Instead of seeing criticism as an obstacle to overcome, I gradually began viewing it as valuable evidence that helped strengthen both my thinking and my research.

Long before I received the degree, I had gradually discovered that professional excellence depends less on the quantity of knowledge accumulated than on how effectively that knowledge is governed. This realisation extended far beyond the specific context of indoor air engineering. I began recognising the same principle whenever I observed experienced professionals making complex decisions, researchers interpreting scientific evidence, educators designing learning experiences, and organisational leaders responding to uncertainty. In every case, success depended not simply on what people knew but on how they governed the development and application of that knowledge within an ever-changing context.

By the time I completed my doctoral research, I no longer viewed Artificial Intelligence, engineering, or even education in quite the same way. I had come to understand that the greatest challenge facing professionals was not acquiring more information but developing the capability to govern the thinking through which information becomes useful judgement and value-oriented action. The doctorate therefore did much more than prepare me to

conduct research. It prepared me to approach every future challenge with greater curiosity, humility, discipline, and responsibility. More importantly, it convinced me that the most enduring contribution I could make throughout the rest of my career would not be helping people accumulate more knowledge, but helping them develop the capability to govern how that knowledge is understood, questioned, refined, and ultimately transformed into meaningful action that creates value for others.

Although I had remained actively involved in the organisation throughout my doctoral studies, completing the Doctor of Engineering fundamentally changed the way I approached my professional responsibilities. The transformation had not occurred overnight or only after I submitted my thesis. It had developed progressively throughout the five years of my research as I continuously refined my understanding, challenged my assumptions, and learnt to govern my own thinking more deliberately. Consequently, by the time I graduated, many colleagues had already begun noticing subtle changes in how I approached complex indoor air investigations, facilitated technical discussions, and supported multidisciplinary teams. The doctorate simply marked the formal recognition of a transformation that had already been taking place in my professional practice.

The organisation also recognised that the impact of my work extended well beyond the successful completion of a doctoral degree. During the five years of my research, the governance-oriented approach had begun improving the transparency, adaptability, and quality of our AI-assisted diagnostic practice while preserving valuable professional knowledge that had previously remained largely within the minds of experienced practitioners. The organisation was becoming increasingly capable of responding to new scientific evidence, refining AI-assisted diagnostic workflows, and strengthening the consistency of professional judgement without relying entirely on lengthy software development cycles.

As these benefits became increasingly evident, senior management expanded my responsibilities and promoted me to lead the organisation's Artificial Intelligence and Digital Innovation initiatives for indoor environmental engineering. The promotion was never presented as a reward for obtaining another academic qualification. Rather, it reflected the organisation's confidence that I could help shape how emerging technologies would be governed and translated into sustainable professional practice that created lasting value for the organisation, its clients, and the communities they served.

My colleagues naturally congratulated me on completing the doctorate, and many were interested in learning about the technical outcomes of my research. However, I quickly realised that the most valuable contribution I could make was not presenting another software solution or introducing another Artificial Intelligence platform. Instead, it was helping my colleagues understand why experienced professionals often arrived at different conclusions despite having access to the same information, and how those differences frequently originated from the way understanding had been developed rather than from the

information itself. I wanted them to recognise that Artificial Intelligence could only be as valuable as the professional reasoning it was designed to represent and support. If professionals did not consciously govern how they developed understanding, interpreted evidence, exercised judgement, and made decisions, even the most sophisticated AI system would merely automate imperfect reasoning.

Before my doctoral journey transformed the way I approached professional practice, I would probably have focused on demonstrating the technical capabilities of the Low-Code AI framework that had emerged from my research. However, my research had convinced me that the technology itself was not the most significant contribution. Its value depended fundamentally on the quality of the professional reasoning it was designed to represent and support. Consequently, I no longer regarded discussions about Artificial Intelligence as discussions that should begin with technology. They had to begin with the thinking that would eventually govern the technology.

Instead, I found myself beginning almost every discussion with questions rather than answers. Before talking about technology, I encouraged colleagues to examine how they interpreted evidence, identified assumptions, evaluated uncertainty, and determined whether their understanding of a situation was sufficiently complete to justify a particular decision. Rather than asking what the AI system should do, I first encouraged them to ask what they themselves should understand before deciding what the AI system ought to represent. Only after we had developed a sufficiently robust understanding of the problem did we discuss how the Low-Code AI framework could faithfully represent, support, and continuously refine that professional reasoning. This simple shift in perspective gradually transformed the nature of our professional conversations and laid the foundation for a different way of practising indoor air engineering.

Initially, this approach surprised many people within the organisation. Some expected the newly graduated Doctor of Engineering to return with sophisticated algorithms, complex automation frameworks, or advanced predictive models that would immediately transform professional practice. I understood their expectations because, before my own transformation, I would probably have expected exactly the same. For many years, I too had believed that technological sophistication naturally represented professional progress. It would have been tempting to demonstrate the elegance of the Low-Code AI framework, the efficiency of the automated workflows, and the impressive capabilities of the Artificial Intelligence models that had emerged from my research. Yet, as I looked around the meeting room and saw experienced colleagues waiting for the technology to provide the answers, I quietly realised that I had once been sitting exactly where they were. The greatest lesson of my doctoral journey was not that Artificial Intelligence had become more capable. It was that I had become more aware of the responsibilities that remained uniquely human.

Although my research had indeed produced technological innovations, I gradually realised that introducing new technology without first strengthening professional reasoning would simply automate existing weaknesses. Artificial Intelligence could process information at remarkable speed, but it could not determine whether the professional understanding embedded within the workflow was contextually appropriate. If the mental model guiding the workflow was incomplete, poorly governed, or based on inappropriate assumptions, the resulting AI-assisted system would simply reproduce those weaknesses more efficiently.

As I reflected upon the investigations that had shaped both my career and my doctoral research, I recognised that the most costly mistakes had rarely arisen because professionals lacked information. More often, they arose because we interpreted available information through incomplete understanding, overlooked important contextual variables, or failed to question assumptions that later proved inappropriate. Those weaknesses had once governed my own thinking as well. The difference was that I could now recognise them more clearly before allowing them to shape my judgement.

Consequently, my early efforts concentrated not on teaching colleagues how to use Low-Code AI, but on helping them recognise the importance of governing the reasoning that would eventually be represented through the technology. Rather than presenting Artificial Intelligence as a replacement for professional judgement, I encouraged colleagues to view it as a powerful extension of professional reasoning that demanded even greater responsibility from those designing and governing it. Gradually, I noticed a subtle but important change in our conversations. Instead of asking how Artificial Intelligence could solve the problem, colleagues increasingly began asking whether their own understanding of the problem was sufficiently complete before deciding what the Artificial Intelligence system should represent. Each time I heard that question, I was reminded of how profoundly my own way of thinking had changed. The transformation I had hoped to achieve through technology was, in reality, beginning with people.

As colleagues gradually embraced this different way of approaching professional practice, the quality of our collaboration began changing in ways that none of us had anticipated. Project meetings no longer revolved primarily around identifying the fastest solution or defending the most convincing explanation. Instead, they increasingly became opportunities for the collective development of understanding. Engineers, consultants, building scientists, software specialists, and other members of multidisciplinary teams contributed their perspectives with a shared recognition that no single individual possessed a complete understanding of the situation. Rather than treating differences in opinion as obstacles to agreement, we began viewing them as opportunities to expose assumptions, identify missing perspectives, and improve the quality of our collective reasoning.

This change was particularly important because the complexity of indoor air investigations rarely allows a single line of evidence to provide a complete explanation. Different stakeholders naturally interpreted the same information through the perspectives of their

own disciplines, experiences, and responsibilities. Instead of attempting to reconcile these differences by selecting one explanation over another, we learnt to examine why different interpretations had emerged in the first place. We explored the assumptions shaping each perspective, identified contextual variables that might have been overlooked, distinguished between evidence and inference, and considered how additional information might strengthen, weaken, or refine our current understanding. As a result, professional discussions became less about reaching immediate agreement and more about progressively developing a richer understanding of the problem before deciding upon the most appropriate course of action.

Looking back, I realised that the greatest cultural change within the organisation was not the adoption of Artificial Intelligence or Low-Code automation. It was the gradual acceptance that sound professional judgement emerges from continuously refining understanding rather than prematurely defending conclusions. As this mindset became increasingly embedded in our daily practice, colleagues developed greater confidence in questioning their own reasoning, seeking alternative perspectives, and revising their judgements whenever new evidence justified doing so. The discussions themselves became more intellectually rigorous, more transparent, and ultimately more productive because our shared objective was no longer to prove who was right but to understand the situation well enough for the organisation to make the most appropriate value-oriented decision.

The practical effects gradually became evident across the organisation. Investigation teams devoted more effort at the outset of projects to understanding the situation comprehensively, yet they spent considerably less time revisiting incorrect diagnoses, repeating site investigations, or revising recommendations after important contextual factors had been overlooked. As a result, investigations became more efficient despite requiring greater reflection at the beginning. Clients also noticed the difference. Our reports no longer presented conclusions as isolated technical findings but explained the reasoning that supported them, the limitations that remained, and the circumstances under which future evidence might justify a different interpretation. Rather than expecting absolute certainty, many clients expressed greater confidence because they could follow the reasoning through which professional judgements had been reached. I gradually realised that professional trust is strengthened not by claiming certainty where none exists but by making professional reasoning sufficiently transparent for others to understand, evaluate, and refine over time.

Another important change emerged in the way professional capability developed within the organisation. Throughout my early career, I had often observed that the greatest asset of experienced practitioners was not merely their technical knowledge but their accumulated judgement. Unfortunately, much of that judgement remained inaccessible to younger engineers because it existed largely as personal experience rather than shared organisational knowledge. Implementing the governance-oriented framework allowed us to

externalise that thinking. Instead of recording only actions, decisions, and technical outputs, our AI-assisted workflows documented the assumptions considered, competing explanations evaluated, contextual factors examined, evidence relied upon, and uncertainties that remained before important decisions were made. Professional expertise therefore became progressively less dependent upon individual memory and increasingly embedded within the organisation itself, allowing valuable knowledge to be preserved, continuously refined, and transferred to future generations of engineers.

Watching younger colleagues grow within this environment became one of the most personally fulfilling experiences of my career. Newly recruited engineers often told me that, for the first time, they understood not only what experienced professionals recommended but why they reached those recommendations. Many remarked that previous training had taught them how to perform engineering tasks, whereas the governance-oriented approach taught them how to think about engineering problems. Their reflections resonated deeply with me because I recognised my younger self in many of them. I had once believed that accumulating knowledge alone would naturally produce sound professional judgement.

My own journey had demonstrated otherwise. Knowledge provides the foundation for professional practice, but judgement depends upon the continual development of understanding, the disciplined interpretation of evidence, the willingness to question assumptions, and the humility to refine one's thinking whenever new information becomes available. Seeing younger professionals develop these capabilities convinced me that the greatest value of my doctoral research lay not in the technology it produced but in the way it helped people become more thoughtful, responsible, and confident professionals.

As these outcomes became increasingly apparent, senior management began viewing the governance-oriented framework as more than a successful research project. It became an important organisational capability. The company was now able to incorporate emerging scientific evidence, revised professional guidance, lessons learnt from completed investigations, and newly acquired expert knowledge into its AI-assisted diagnostic workflows far more rapidly than had previously been possible. This significantly reduced dependence on lengthy software redevelopment while enabling experienced indoor air professionals to govern the continuous evolution of the system themselves. The organisation became more adaptive, more resilient, and better prepared to respond to future challenges because valuable professional knowledge was no longer confined to individual experts but had become part of a continuously evolving organisational asset.

Looking back, I realised that the most enduring contribution of my doctoral research was not simply making Artificial Intelligence more accessible through Low-Code platforms. It was demonstrating that sustainable innovation depends upon enabling professionals to govern the continual evolution of their own reasoning. That insight would eventually shape not only the future direction of my professional career but also my subsequent work as a scholar, educator, and mentor.

As the organisational benefits became increasingly visible, interest in the work gradually extended beyond the company. Professional institutions, engineering societies, universities, government agencies, and industry organisations began inviting me to share our experiences implementing the governance-oriented framework in authentic engineering practice. What initially started as technical presentations on Low-Code AI and AI-assisted diagnostic workflows soon evolved into broader discussions about organisational learning, professional capability development, and the future relationship between engineering expertise and Artificial Intelligence. These conversations frequently continued long after the formal presentations had ended because participants recognised that the issues extended far beyond software or automation. They reflected challenges that many organisations were already facing as they sought to integrate increasingly capable Artificial Intelligence into professional practice without diminishing the quality of human judgement or organisational knowledge.

These engagements also provided me with an invaluable opportunity to learn from professionals working in different industries, organisational settings, and cultural contexts. Rather than simply presenting the outcomes of my own research, I listened carefully to their experiences, the challenges they encountered, and the questions that remained unresolved within their own organisations. Although the technical applications differed considerably, many described remarkably similar concerns regarding the preservation of professional expertise, the transparency of AI-assisted decisions, the continuous adaptation of knowledge, and the difficulty of ensuring that technology evolved alongside professional practice. Those conversations reinforced my belief that the principles developed during my doctoral research addressed a challenge that extended well beyond indoor air engineering. More importantly, they reminded me that meaningful scholarship should continue learning from practice rather than merely seeking to influence it.

These engagements gradually evolved beyond professional development activities into long-term scholarly collaborations. Universities, research institutes, government agencies, and professional organisations became interested in exploring whether the governance-oriented design theory could contribute to challenges extending well beyond indoor air diagnostics. What I found particularly rewarding was that these collaborations rarely began with discussions about my previous work. Instead, they began with authentic problems that researchers and practitioners were already struggling to solve within their own disciplines. This created opportunities for us to examine together whether the underlying principles could illuminate new contexts rather than simply reproduce earlier successes. Every new collaboration exposed me to different ways of thinking, different professional cultures, and different forms of evidence, all of which broadened my own understanding far beyond the boundaries of my original field.

I welcomed these opportunities because they encouraged rigorous dialogue rather than uncritical acceptance. Some collaborators questioned the scope of the theory, others challenged its assumptions, while many proposed extensions that I had never previously considered. Far from seeing these discussions as criticisms to be resisted, I regarded them as an essential part of responsible scholarship. Each collaboration became an opportunity to examine where the design theory was robust, where it required adaptation, and where its limitations became apparent under different organisational, cultural, or disciplinary conditions. Those experiences reminded me that no theory remains valuable simply because it has once been successful. Its continued relevance depends upon its willingness to evolve alongside new evidence, new contexts, and new questions. My doctoral journey had fundamentally changed my relationship with knowledge itself. I no longer regarded scholarship as the pursuit of being right but as the continual pursuit of becoming less wrong through disciplined inquiry, open dialogue, and thoughtful reflection.

As my publications gradually accumulated, I became increasingly aware that the scholarly community was recognising something quite different from what I had originally imagined. When I first embarked upon my doctoral research, I had assumed that the principal contribution would be the Low-Code AI framework itself. Over time, however, it became evident that what attracted sustained academic interest was the conceptual foundation that explained why the framework had been designed in that particular way. Rather than being recognised for another Artificial Intelligence application, I found myself contributing to a broader conversation about how professionals and intelligent technologies should evolve together.

This perspective gradually encouraged new lines of inquiry connecting engineering design, cognitive governance, Artificial Intelligence, educational theory, organisational learning, and professional decision-making. Perhaps the most humbling experience was seeing researchers whom I had never met interpret, challenge, extend, and adapt these ideas within disciplines that I had never anticipated influencing. Although their applications differed substantially from my own, they all shared a common commitment to strengthening the quality of human judgement in increasingly technology-enabled environments. Watching those ideas develop independently of my own work was deeply satisfying because it demonstrated that they had become part of a much larger scholarly conversation.

Reflecting upon that period, I realised that research achieves its greatest significance not when it remains associated with its original author but when it enables others to ask better questions, develop new knowledge, and create value in ways that could never have been predicted at the outset. That realisation brought me a profound sense of fulfilment because it reminded me that the purpose of scholarship is not to preserve one's own ideas unchanged but to create foundations upon which future generations can continue building.

Seventeen years after completing my Doctor of Engineering, I accepted an appointment as Professor of Engineering Design and Professional Practice while continuing to serve in a strategic leadership role within the organisation where my doctoral journey had begun. The appointment recognised that my scholarship had been shaped by sustained engagement with real engineering challenges. At the same time, the organisation continued to benefit from the knowledge, educational innovation, and research generated through my academic work. Rather than asking me to choose between academia and industry, both institutions deliberately created an arrangement that enabled each to strengthen the other in pursuit of advancing engineering knowledge, professional capability, and societal value.

Some colleagues wondered whether it would be possible to contribute meaningfully to both environments without compromising either. I understood their concern because many professionals eventually commit exclusively to academia or industry as their careers progress. My own experience, however, had led me to a different conclusion. I no longer regarded teaching, research, and professional practice as separate careers but as interconnected responsibilities serving a common purpose.

Professional practice continually revealed authentic challenges requiring deeper investigation. Research transformed those challenges into new knowledge, while teaching prepared future professionals to apply, evaluate, and extend that knowledge within their own contexts. As that knowledge was applied in professional practice, it generated fresh insights, new evidence, and previously unrecognised questions, which in turn became the starting point for further scholarly enquiry and educational innovation. Together, these activities formed a continuous cycle of learning, application, reflection, and improvement that reflected the principles of cognitive governance which had become the foundation of everything I sought to achieve through research, education, and professional practice.

My philosophy as an educator differed substantially from many of the approaches I had experienced during my own education. Although technical knowledge remained indispensable, I regarded it as only one element of professional capability. My primary responsibility was not simply to teach engineering concepts but to help learners develop the capacity to interpret complex situations, exercise sound judgement, and make value-oriented decisions. Consequently, every lecture, design studio, research supervision, and industry workshop began by developing an understanding of the problem before considering possible solutions.

Rather than introducing software, analytical methods, engineering standards, or Artificial Intelligence at the outset, I first encouraged learners to construct a mental model of the situation they were seeking to understand. Together, we explored the stakeholders involved, their objectives, the contextual factors influencing the problem, the assumptions shaping current interpretations, the uncertainties that remained, the potential consequences of different decisions, and the relationships connecting these various elements.

Only after this foundation had been established did we examine how engineering knowledge, Artificial Intelligence, digital technologies, and computational tools could support, refine, and extend their reasoning. This sequence proved remarkably important because it ensured that technology was introduced as a means of strengthening professional judgement rather than directing it. By establishing understanding before selecting tools, learners became better equipped to evaluate whether a technological solution was appropriate, how it should be applied, and when human judgement ought to prevail. Over time, I observed that this approach not only enhanced learners' technical competence but also strengthened their confidence, intellectual independence, and ability to navigate unfamiliar problems with thoughtful and responsible judgement.

One of the most rewarding aspects of teaching was witnessing the gradual evolution of my students. During the early weeks of many courses, classroom discussions typically centred on how to use the prescribed software effectively, how to obtain the correct answer, or how to construct prompts that would produce the desired output from Artificial Intelligence. These enquiries reflected a natural tendency to focus on tools and outcomes before fully appreciating the importance of the thinking that should guide their use. As the semester progressed, however, I began noticing a subtle but profound change. Students increasingly examined whether they had explored competing explanations fairly, recognised the assumptions influencing their conclusions, considered relevant contextual factors, and gathered sufficient evidence to justify their decisions. Rather than seeking immediate answers, they became more comfortable engaging with uncertainty and refining their perspectives before deciding how best to proceed.

Observing this transformation brought me immense satisfaction because it reflected a capability that extended far beyond technical proficiency. I could see students becoming more reflective, intellectually curious, and increasingly confident in navigating complex situations. They were learning to pause, evaluate different perspectives, interpret evidence carefully, and exercise sound judgement before taking action. Whenever I witnessed this shift, I knew that the most meaningful development had already occurred, long before it could be demonstrated through a final assessment or examination. In those moments, I was reminded that the enduring purpose of education is not merely to transfer knowledge, but to cultivate the wisdom, responsibility, and discernment required to apply it in ways that create lasting value for others.

Supervising postgraduate researchers became another deeply fulfilling dimension of my academic career. Rather than encouraging students to defend predetermined hypotheses or pursue fashionable technologies, I guided them to view scholarly enquiry as a disciplined process of progressively refining their interpretation of complex problems. We often reminded ourselves that our objective was not to confirm our initial expectations but to discover what the evidence genuinely revealed, even when it challenged our assumptions. This perspective cultivated intellectual honesty, encouraged openness to

alternative explanations, and reduced the temptation to become personally attached to particular conclusions.

One of the greatest satisfactions of academic supervision came from observing how students gradually matured as independent scholars. Many later reflected that the most valuable lesson they had gained was not a specific methodology, analytical technique, or technological skill, but the discipline of critically examining their assumptions, evaluating evidence objectively, and remaining willing to revise their interpretations whenever new insights emerged. Those reflections reassured me that my contribution extended beyond helping students complete successful theses. I was witnessing the development of scholars capable of approaching uncertainty with intellectual humility, disciplined judgement, and a genuine commitment to advancing knowledge in service of society.

As the years passed, opportunities to contribute at national and international levels became an increasingly natural extension of my work. I served on advisory panels, participated in developing professional guidelines, collaborated with interdisciplinary research teams, and contributed to educational initiatives relating to Artificial Intelligence, engineering education, and professional practice. These engagements allowed me to work alongside scholars, practitioners, policymakers, and industry leaders who shared a common commitment to addressing complex societal challenges through responsible innovation. Each collaboration broadened my perspective by exposing me to different organisational contexts, cultural viewpoints, and disciplinary traditions. In many respects, I learnt as much from these engagements as I contributed to them.

Although I was deeply honoured by the trust placed in me, I remained mindful that recognition carries responsibility rather than entitlement. Every invitation provided an opportunity to listen, to learn, and to contribute to conversations that were far larger than my own work. Whenever I accepted an award or delivered a keynote lecture, I reminded myself that the occasion was not about celebrating personal achievement but about advancing ideas that could help others address challenges within their own professional contexts. The true value of scholarship lies not in the visibility of the scholar but in the lasting influence of the knowledge shared with others.

Looking back, I gradually realised that the distinctions between engineer, researcher, educator, and professional leader had become increasingly indistinct. Rather than representing separate identities, they had evolved into different expressions of the same lifelong commitment. Whether I was investigating a complex indoor air problem, supervising a doctoral student, designing an educational programme, developing a Low-Code AI workflow, reviewing a research manuscript, or advising policymakers, my responsibility remained fundamentally unchanged: to help people develop the capability to interpret evidence thoughtfully, exercise sound judgement, and act responsibly in situations characterised by complexity and uncertainty.

That purpose gave coherence to every role I undertook and reminded me that meaningful influence is rarely defined by positions held, honours received, or technologies created. Instead, it is reflected in the extent to which our work enables others to think more clearly, make wiser decisions, and create enduring value within their own communities. As I reflected upon my journey, I realised that if I had left behind anything of lasting significance, it would not simply be a design theory or an AI framework, but a way of thinking that empowered others to govern themselves before attempting to govern the increasingly intelligent technologies shaping our world. **The End!**