18-PART SERIES
HEAT PUMPS: YOUR BURNING QUESTIONS, ANSWERED NOW

10/18

Heat Pumps and AI – a Perfect Match?

Author: Dr.-Ing. Marek Miara, published: 19.01.2026

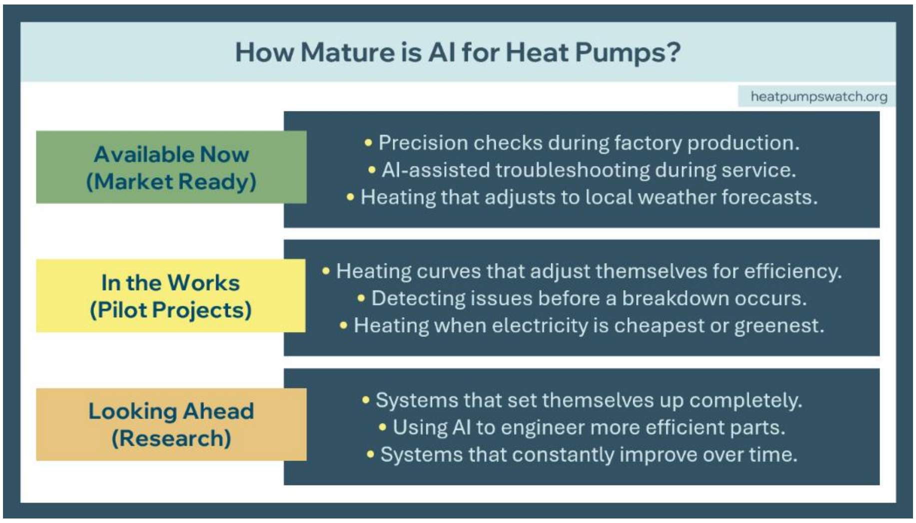

Old Necessity, New Tools

Warmth in winter, coolness in summer—these needs have accompanied humanity for thousands of years. Heat pumps fulfill them with sophisticated technology. Artificial intelligence, on the other hand, is young and still experimental in many areas. Can both worlds come together in a meaningful way?

The answer requires differentiation. From development to operation, there are various steps: design, production, planning, installation, commissioning, operation, and maintenance. At each of these points, there is potential for optimization and also for the use of different forms of artificial intelligence. However, “AI” is often confused with other principles or used as a *buzzword* for marketing purposes. This article examines where real machine learning is used – and shows the varying market readiness of the different applications.

What “Intelligence” Actually Means

True artificial intelligence is characterized by three properties: It learns from experience. It recognizes patterns that no one has explicitly programmed. And it adapts to new situations.

However, most systems advertised as "intelligent" are rule-based: they follow hard-coded instructions—such as threshold-based controls that raise the flow temperature by a fixed value at certain outside temperatures. Such parametric controls are automation, not machine learning.

The key difference: rule-based systems operate according to predefined logic. Learning systems recognize correlations from data and adapt their behavior without explicit programming.

The Value-added Chain: From Design to Service

The path from development to ongoing operation of a heat pump involves several phases: product design and production at the manufacturer, planning and dimensioning for a specific building, on-site installation, commissioning and parameterization, operation over many years, and regular maintenance. Each of these steps offers specific opportunities for optimization—and different ways in which machine learning can contribute.

The following sections systematically examine this chain. For each step, we analyze which AI applications already exist, which are being developed in research, and where the term is merely used for marketing purposes.

1. Desing and Production

The manufacture of a heat pump involves numerous steps: from the development of optimized components to the production of individual parts, assembly, and final inspection. Each of these steps offers potential for optimization—whether through more efficient designs, more precise manufacturing, or reliable quality assurance.

Product development is all about getting the most out of the components. The search for the perfect shape is particularly important for heat exchangers and compressors: they should transfer as much heat as possible while offering minimal resistance to the flow of air or water. Only then can the heat pump operate at maximum efficiency and consume little electricity. To achieve this ideal, engineers must test countless combinations of shapes and angles. Traditionally, they rely on years of experience and test different designs step by step on the computer until a satisfactory result is achieved.

AI can expand the experience and intuition of engineers and find new, more efficient designs.

It is being used more and more.

Machine learning expands this approach: algorithms evaluate millions of simulations and suggest geometries that lie beyond human intuition. One example is the application of symbolic regression (automatically finds an equation that best describes given data) to optimize turbo compressors, achieving energy savings of 20–25%[1]. Such methods are increasingly finding their way into research and development.

In manufacturing, the focus is on quality control. Hairline cracks in solder joints, deviations in sensor positioning, or irregularities in heat exchangers are difficult to detect visually but can affect service life. Image recognition systems, trained on thousands of labeled examples, achieve detection rates that surpass those of human inspectors. These applications are state of the art in industrial manufacturing.

2. Planning

Before a heat pump is installed, it must be dimensioned. How much heating capacity does this specific building need? The answer determines the success or failure of the entire system.

Field studies show (see Episode 2) that oversizing heat pumps is a common problem. If the installed capacity significantly exceeds actual demand, the system often operates in partial load mode or cycles on and off. This can lead to suboptimal efficiency and increased wear and tear.

AI can assist with heat pump sizing.

Normal algorithms are often used, not "true" AI.

There are many reasons for this: Standard heating load calculations are based on design conditions that only occur for a few hours a year. At the same time, actual user behavior is difficult to predict. The result is a discrepancy between calculated and actual heat demand.

Data-based approaches could remedy this situation. Systems that learn from the operating data of thousands of installed systems could make realistic predictions: Which heat pump configuration has actually achieved the best performance in comparable buildings? Such empirical models would supplement normative calculations, not replace them.

What is sold as AI but, in fact, is not:

Most "AI configurators" on manufacturer websites are something else entirely. They are digital selection aids based on standards and databases. Useful, time-saving—but not learning. The user provides the data, the system determines the appropriate information, and the process is complete. This is a digital product catalog with a convenience feature, not artificial intelligence.

3. Installation

Errors still occur during the installation phase, which impair subsequent efficiency and cause unnecessary costs.

Genuine AI applications are rare at this stage. Most "smart" installation aids are limited to digital checklists or augmented reality instructions (where digital information is superimposed on the field of vision) – in most cases, these are rule-based systems without learning capabilities.

AI has played a minor role in installation so far.

Research projects are investigating self-learning approaches: The Danish CEDAR project is developing systems that use reinforcement learning (where the system finds optimal strategies through trial and evaluation) and digital twins to perform hydraulic balancing automatically – an "install-and-forget" concept for retrofit applications[2]. The Swiss University of Applied Sciences ZHAW is working on AI-supported commissioning automation[3].

Currently, the greatest lever for quality continues to be professional execution by trained personnel.

4. Bringing into Service / Starting

After installation, the heat pump must be adjusted. The so-called heating curve determines how hot the water flowing through the heating system should be, depending on the outside temperature. This setting is made once and then usually never touched again.

The problem is that buildings change. New windows affect heat requirements, new residents bring different usage habits, and building fabric ages. Once set, the heating curve does not automatically adapt to these changes and may become suboptimal after a few years.

During commissioning, AI provides assistance that would otherwise be difficult to obtain. The control system can be customized to the individual based on various parameters.

In this case, research shows that “real” AI can help. A neural network observes the building over several days. It measures how quickly rooms cool down, how much the sun heats them up, and when they are heated. From this data, it learns a model of thermal behavior—without explicit programming of the building physics.

With this learned model, the system can regulate proactively: solar gains in the afternoon are anticipated, recurring usage patterns are recognized, and the heating is adjusted accordingly.

In simulations and field tests, such systems achieve efficiency gains of 5–13% compared to static heating curves, while also improving comfort. The Fraunhofer Institute for Solar Energy Systems has developed and tested such adaptive controllers in collaboration with partners from industry and research as part of the AI4HP project[4].

However, many systems described as “self-learning” work with simple pattern recognition: the system stores when the user manually changes the temperature and repeats this time program. This corresponds to the functionality of programmable thermostats and is not machine learning.

5. Operation Phase

There is an approach called “model predictive control.” The system creates a mathematical model of the building and calculates in advance how changes in the heating system will affect the temperature. This allows it to optimize operating modes, take electricity prices into account, and maximize the building's own consumption of solar power.

According to studies, this control system enables cost savings of up to 40% through optimized operational management. The heat pump runs preferentially when electricity prices are low, uses higher storage temperatures for temporal flexibility, and also improves grid support by 13% – despite a slightly lower COP[5].

However, we must be honest here: this is not artificial intelligence in the true sense of the word. It is mathematical optimization. The system does not “learn” from experience—it calculates the optimum based on equations defined by an engineer. The approach dates back to the 1980s and has proven itself since then.

That doesn't detract from its usefulness. But if you want to understand where real AI makes a difference, you need to be aware of this limitation. Optimization calculates the best outcome based on known rules. By definition, AI is supposed to discover rules that no one knew before.

The more time passes, the more effectively and efficiently the AI can monitor the operation of the heat pump.

Where true AI works during the operation phase:

Reinforcement learning takes a different approach. The system tries out different strategies and evaluates the results. Was the temperature comfortable? Was power consumption low? Over time, it learns which actions produce the best results in which situations—without anyone having to specify the rules.

A concrete example is defrost control. In air source heat pumps, ice forms on the evaporator at low temperatures and must be defrosted regularly. Conventional systems usually do this at fixed intervals or at certain temperatures. A learning system recognizes the patterns in the data and finds the optimal time itself – not too early (energy waste), not too late (performance drop).

6. Servicing

When a heat pump is not running optimally, you often only notice it when the electricity bill arrives – or when the system fails completely. Creeping problems such as slow refrigerant loss, a dirty heat exchanger, or a worn valve often go undetected for a long time.

A system that knows how a heat pump normally operates can detect deviations. It knows how temperatures and pressures normally relate to each other and how much power the compressor draws under certain conditions. If these patterns change gradually, this indicates a problem – often weeks or months before a limit value is exceeded or an error code appears.

This is fundamentally different from traditional monitoring. A conventional control system reports: “Error E17: Low pressure fault” – when it is too late. A learning system could say: “Efficiency has dropped by 3% in the last four weeks. Probable cause: evaporator contamination. Recommendation: cleaning next month.”

Research has made Big Leaps Forward:

Systematic reviews of the application of AI in building technology, which evaluated 230 studies, show that fault diagnosis is the most technologically mature area of application[6]. Deep learning-based systems for automatic fault diagnosis in HVAC systems achieve accuracies of over 97%[7]. When detecting refrigerant leaks, for example, the best algorithms achieve hit rates of over 95%.

AI can detect such problems in real time and with little effort, problems that would otherwise only become apparent much later.

Particularly promising are so-called “soft faults” – problems that do not cause acute malfunctions but reduce efficiency. A slightly dirty condenser, a stiff valve, an aging compressor. Detecting these problems before they lead to failure is the real added value.

Predicting when a component is likely to fail goes one step further. Certain components—for example, the electrolytic capacitor in the inverter or the reversing valve—show characteristic wear patterns. Recognizing these patterns allows maintenance to be planned: not too early (waste of money), not too late (unplanned failure).

Such systems are well established in industry. For heat pumps in residential applications, they are still in development—but the potential is enormous.

7. Grid Integration

Heat pumps could play an important role in the power grid. They can store heat—in the building itself, in the hot water tank, in the buffer storage tank. This makes them flexible consumers that can balance out peak loads.

The British LATENT project at the University of Southampton[8] shows that AI-controlled heat pumps can reduce their load by up to 90% during peak times without compromising on comfort. The system preheats when electricity is cheap and plentiful and pauses when the grid is under pressure.

AI can monitor electricity prices and regulate the operation of the heat pump accordingly in a cost-efficient manner.

However, this only works with predictive control. The system needs to know what the weather will be like, when residents will come home, and how electricity prices will develop. This is where AI can make a real contribution.

6 Conclusion

Established heat pump technology can certainly benefit from artificial intelligence.

Manufacturers can use AI-supported design to develop more efficient components and, for example, increase production quality using computer vision. Planners benefit from data-based design tools based on empirical operating data. Installers could use self-learning commissioning systems in the future. End users gain efficiency and comfort with reduced operating costs thanks to adaptive control. The energy system benefits from intelligent load shifting, which better balances supply and demand.

It is crucial to distinguish between true machine learning and rule-based systems that are simply referred to as “AI.” True AI is characterized by the ability to learn from data, adapt without explicit programming, and behave predictively. Many systems advertised as intelligent do not meet these criteria—a confusion that is often exploited for marketing purposes.

Research into intelligent heat pump controls is highly advanced, but widespread market introduction is still lacking. Validated efficiency gains result from adaptive control, significant cost savings through load shifting, and high hit rates in machine learning for troubleshooting. The potential is demonstrable – but widespread practical application will take time.

[1] Shalash, K., and Schiffmann, J. (February 3, 2020). "Pressure Profile Measurements Within the Gas Film of Journal Foil Bearings Using an Instrumented Rotor With Telemetry." ASME. *J. Eng. Gas Turbines Power*. March 2020; 142(3): 031013. https://doi.org/10.1115/1.4044798
[2] Poulsen, J. L., Aguilera, J. J., Madsen, H., & Markussen, W. B. (2023). *Country Summary for Denmark on Digitalization and IoT for Heat Pumps*. IEA Heat Pumping Technologies Programme, Annex 56 – Digitalization and IoT for Heat Pumps. https://heatpumpingtechnologies.org/annex56/
[3] ZHAW Zurich University of Applied Sciences (2023). *Automated Commissioning of Heat Pumps*. Institute of Applied Mathematics and Physics (IAMP), ZHAW. Project completed 02/2023–05/2023. https://www.zhaw.ch/en/research/project/73687.
[4] Fraunhofer Institute for Solar Energy Systems (ISE) (17 Dec 2024). *AI-Controlled Heat Pumps Increase Efficiency*. Press release. https://www.ise.fraunhofer.de/en/press-media/press-releases/2024/ai-controlled-heat-pumps-increase-efficiency.html.
[5] Tomás, A., Lämmle, C., & Pfafferott, J. (2025). Demonstration and Evaluation of Model Predictive Control (MPC) for a Real-World Heat Pump System in a Commercial Low-Energy Building for Cost Reduction and Enhanced Grid Support. *Energies*, *18*(6), 1434. https://doi.org/10.3390/en18061434
[6] Boutabba, A., Albalkhy, W., Lafhaj, Z., Roussel, J., Yim, P., Danel, T. (2025): Enhancing Thermal Comfort and Energy Efficiency in Buildings Using Artificial Intelligence: A Systematic Literature Review. Modular and offsite Construction Summit, Montreal, Kanada.
[7] Zhang, F., Saeed, N., & Sadeghian, P. (2023). *Deep learning in fault detection and diagnosis of building HVAC systems: A systematic review with meta analysis.* Energy and AI, 12, 100235. https://doi.org/10.1016/j.egyai.2023.100235
[8] University of Southampton (2025). *LATENT: Residential Heat as an Energy System Service*. EPSRC-funded research project on residential heating system flexibility and demand control, University of Southampton. Available at: https://www.southampton.ac.uk/research/projects/latent-residential-heat-as-an-energy-system-service (accessed Jan 2026).