

Aalto University

MS-E2177 - Seminar on Case Studies in Operations Research D

Artificial Intelligence in Water Network Planning

Team

Bilqays Ayoub

Emma Kämäräinen

Leevi Pankkonen (Project Manager)

Sini Poikonen

Client: Sweco

May 29, 2026

Contents

1	Introduction	4
	Background	4
	Objectives	4
2	Literature Review	5
	What is Artificial Intelligence?	5
	Machine Learning and Deep Learning	6
	Subtypes of Machine Learning	6
	Neural Networks and the Deep Learning Process	7
	Generative AI and Large Language Models	7
	Categories of Generative AI	7
	Transformer Architecture and GPT Models	8
	Classification Frameworks for AI Systems	8
	OECD Framework	9
	CSET Autonomy Classification	9
	NIST AI Use Taxonomy	10
	Summary	11
	Frameworks	11
	Theory of Change: Logic and Signal and Signal Based Strategy	11
	Automation–Augmentation–Autonomy	12
3	Data & Methods	13
	Expert interviews	13
	Research Objectives	13
	Methodological Frameworks	13
	Software Interviews	14
	Research Objectives	14
	Workshop	14
	Research objectives	15
4	Results	15
	AI in Water Infrastructure Planning Based on Literature	15
	The Nature of Water Infrastructure Design	15
	AI Opportunities in Water Infrastructure Planning	16
	Limitations and Constraints	16
	Expert interview results	17
	Potential AI Applications in the Planning Process	17
	Expert Insights on Responsibility and the “Black Box” Problem	18
	Barriers to Adoption and Organizational Change	18
	Software Vendor Interview Results	18
	Planning process in general water network	19
	Mock-up: AI-Assisted Source Data Synthesis	20

Functional Overview and Technical Implementation	20
Data Integration and Sources	20
Critical Risks and Technical Uncertainties	21
Scenarios	21
Workshop Results	22
5 Discussion	22
6 Conclusions	23
7 Self Assessment	26
Following the Plan	26
Project Success	26
Failures	26
Improvements	27

1 Introduction

Background

Current design workflows in civil water network planning are heavily based on expert-driven manual work. Despite strong professional know-how, many tasks in water supply planning contain substantial routine elements, standardized steps, and extensive data processing. These characteristics suggest potential for applying artificial intelligence (AI) to support the overall planning process.

This project examines how AI could help in water supply planning, both in general and specifically within Sweco's current practices. The research question addresses how AI could affect water supply planning from Sweco's operational perspective. The approach combines management consulting frameworks with operations research methodologies to systematically examine water supply planning processes, focusing on the Finnish water sector.

The project is structured around five work packages: high-level assessment of AI tools, mapping of Sweco's planning processes, review of AI solutions, assessment of organizational change and impact, and development of recommendations with an implementation roadmap. Data collection is based on interviews and workshops with Sweco personnel, combined with desktop studies and literature review.

Objectives

The primary objective is to improve the efficiency of water supply planning by identifying the most suitable AI use cases. The study clarifies what AI can actually do compared to traditional automation, helping to understand where AI adds real value to the planning process.

The project maps Sweco's existing workflows, data flows, and deliverables to understand current practices and identify improvement areas. This mapping helps evaluate different AI solutions (such as language model tools, but also more traditional machine learning methods) in terms of how well they fit Sweco's operations.

The study examines how AI adoption would affect work processes, including changes to daily tasks, skill requirements, and quality control. Finally, the project develops practical recommendations and an implementation roadmap for Sweco, covering both short-term quick wins and long-term strategic changes in technology, skills, processes, and organizational structure.

2 Literature Review

What is Artificial Intelligence?

Artificial intelligence (AI) refers to software and systems capable of interpreting data, making decisions, and performing tasks that resemble human reasoning, without relying on fully pre-defined rules (Delipetrev et al., 2020; Bory et al., 2024). While AI systems may make use of explicit rules, what distinguishes them is their capacity to learn from data and adapt based on experience. This sets them apart from conventional software automation, which requires exhaustive rule-sets to function correctly. For instance, spreadsheet tools where each cell formula must be defined individually (Tatipamula, 2025).

Traditional rule-based automation systems are effective for tasks with a predictable, consistent input and context, but they struggle when inputs change, outputs are ambiguous, or tasks require nuanced judgment. By contrast, AI systems leverage probabilistic reasoning, pattern recognition, and contextual understanding. As an illustrative example, rule-based building management systems have been found to operate roughly 15–20% less efficiently than adaptive AI-driven systems that learn from environmental conditions and usage patterns (Tatipamula, 2025).

It is important to note, however, that while some AI systems can learn from small numbers of examples, the majority of widely used models require large volumes of data to function reliably.

AI research traces its origins to the 1950s. Alan Turing introduced what became known as the Turing Test: the idea that a computer could be considered intelligent if it could sustain a conversation with a human without the human realising they were talking to a machine (Delipetrev et al., 2020).

In the following decades, so-called symbolic AI dominated the field. These systems operated on pre-defined rules and logical facts. Although powerful within narrow domains, symbolic systems proved too rigid to represent the complexity and variability of the real world (Delipetrev et al., 2020; Bory et al., 2024).

From the 1990s onward, machine learning and deep learning emerged as the dominant paradigms. By the 2010s, AI was being used in medical diagnostics, search engines, and recommendation algorithms. For example, Spotify and Instagram rely on machine learning to personalise content. In the early 2020s, deep learning surpassed other methods in performance across a wide range of tasks. A landmark moment came in 2022, when OpenAI released ChatGPT, a large language model based on deep learning that requires vast quantities of training data (Delipetrev et al., 2020; He et al., 2025; IBM, 2024a).

It is useful to distinguish between three conceptual categories of AI (Bory et al., 2024):

Narrow (weak) AI is designed to perform specific, well-defined tasks. It is fast and effective within its domain but cannot generalise or transfer knowledge across different task areas. All currently deployed AI systems fall into this category.

General (strong) AI refers to a hypothetical system with human-like intelligence, capable of performing any intellectual task a human can. Strong AI remains a theoretical concept and has not yet been realised.

Superintelligent AI is a further hypothetical form of intelligence that would surpass human cognitive capabilities across virtually all domains, including creativity and problem-solving. This concept belongs to the realm of science fiction and ongoing theoretical debate, and it raises significant concerns among researchers.

AI is advancing rapidly, and the most capable current systems may appear to approach strong AI, but they are still classified as narrow AI.

Machine Learning and Deep Learning

Subtypes of Machine Learning

Machine learning is the subfield of AI in which systems learn from data rather than from explicitly programmed rules. Its main paradigms differ in how learning occurs (Delipetrev et al., 2020; IBM, 2024a):

Supervised learning involves algorithms that learn to map inputs to outputs based on labelled training examples prepared by humans. A classic illustration is training a classifier to recognise cats in photographs by showing it many labelled images. Supervised learning requires substantial labelled data.

Unsupervised learning involves algorithms that identify previously unknown structures in data by forming clusters, without requiring labelled examples. Given a collection of images containing both cats and dogs, an unsupervised system might discover and group the two categories without being told which is which.

Reinforcement learning studies how agents learn autonomously. Agents receive positive or negative feedback based on how they interact with an environment, and learn from this feedback rather than from labelled data. The AlphaGo system, which learned to play the board game Go at superhuman level by playing against itself, is a well-known example.

Deep learning builds on multi-layer neural network architectures loosely in-

spired by the structure of the human brain. It underlies modern applications such as voice assistants (Siri), chatbots, and large language models including ChatGPT.

Neural Networks and the Deep Learning Process

A neural network is a computational model consisting of layers of interconnected nodes (neurons) that process and transform input signals (Goodfellow et al., 2016; Rumelhart et al., 1986). Its main components are:

- Input layer: receives the raw data features.
- Hidden layers: learn intermediate representations of the data.
- Output layer: produces predictions or classifications.
- Weights and biases: adjustable parameters optimised during training.

Each artificial neuron computes a weighted sum of its inputs, adds a bias term, and passes the result through a non-linear activation function σ :

$$\text{output} = \sigma \left(\sum_i w_i x_i + b \right). \quad (1)$$

Training proceeds in four steps. In the forward pass, the input propagates through the network layer by layer, with each neuron computing its weighted activation. A loss function (e.g. cross-entropy or mean squared error) then measures the discrepancy between the network's predictions and the correct labels. Backpropagation computes the gradient of the loss with respect to each weight using the chain rule, propagating error signals backwards through the network. Finally, an optimisation algorithm updates the weights in the direction that reduces the loss (Goodfellow et al., 2016; Rumelhart et al., 1986).

Generative AI and Large Language Models

Categories of Generative AI

Generative AI (GAI) refers to systems that produce new content—text, images, code, or other outputs, rather than simply classifying or retrieving existing information. Three broad categories can be distinguished (He et al., 2025; IBM, 2024a):

Rule-based GAI generates outputs according to explicit, hand-crafted rules supplied to the system.

Model-based GAI learns from examples. Statistical generative models make inferences and interpret data, while graphics-based models synthesise visual content through physical simulation.

Deep learning-based GAI is built on neural network architectures that imitate the structure of biological neural networks. This category includes the most capable and widely used contemporary systems.

Transformer Architecture and GPT Models

The transformer architecture, introduced by Vaswani et al. (2017), replaced earlier recurrent neural networks (RNNs) with a self-attention mechanism that enables parallel computation over entire input sequences. This made training substantially faster and allowed scaling to models with billions of parameters.

Generative Pre-trained Transformers (GPT) use the decoder component of the transformer. They are pre-trained on large text corpora using next-word prediction and subsequently fine-tuned for specific tasks. Three key properties distinguish transformer-based language models:

1. Scale: more data and more parameters unlock qualitatively new capabilities.
2. Self-supervision: pre-training on unlabelled text enables broad generalisation.
3. RLHF (Reinforcement Learning from Human Feedback): aligns model outputs with human preferences during fine-tuning.

Large language models (LLMs) such as ChatGPT and Claude are text-focused generative tools built on this architecture. Their behaviour is highly sensitive to the prompt provided by the user: the phrasing and specificity of an instruction substantially affect output quality (NVIDIA, 2026; IBM, 2024b; MLOps Community, 2026). These models can process and generate not only text but also images and structured files (e.g. spreadsheets or word-processing documents). They may be deployed as general-purpose tools or fine-tuned for enterprise use. For example, a company-specific deployment in which inputs are not used to retrain the underlying model. A notable concern is energy consumption: LLM inference requires considerable computational resources (Kenny et al., 2025; NVIDIA, 2026).

Classification Frameworks for AI Systems

Several complementary frameworks exist for classifying and evaluating AI systems. This section covers three that are particularly relevant to engineering design contexts.

OECD Framework

The OECD Framework for the Classification of AI Systems (OECD, 2022) provides a contextual perspective structured around four interacting dimensions, with people and the planet at the centre:

- Economic context: industrial sector, business function, criticality, scale, and development stage.
- Data and input: origin, collection method, dynamism, structure, format, rights, anonymisation, and quality.
- AI model: architecture (machine learning, symbolic, hybrid), interpretability, and deployment.
- Task and output: the system’s function (recognition, personalisation, etc.), its level of autonomy, and its principal application domains.

The framework’s primary purpose in this project is to describe how an AI system is used in a particular project and context, rather than to characterise it solely by technical properties. It is particularly well suited to describing LLMs, which resist classification by rigid taxonomies: the same model may represent a low-risk or high-risk system depending entirely on how and where it is deployed.

CSET Autonomy Classification

The Center for Security and Emerging Technology (CSET) approach (Aiken, 2021) classifies AI systems according to observable system properties—most importantly, the degree of autonomy with which the system processes inputs, produces decisions, and acts without human involvement. Three levels of autonomy are distinguished:

- Action autonomy: the system produces a decision *and* executes the action independently.
- Decision autonomy: the system produces a decision output, but a human carries out the action.
- Detection autonomy: the system identifies and surfaces information, but a human makes the final judgement and decision.

CSET’s autonomy-based classification is better suited than the OECD framework for characterising AI tools used within engineering design workflows, where the degree of human oversight is a critical variable.

NIST AI Use Taxonomy

The NIST AI Use Taxonomy (National Institute of Standards and Technology, 2024) complements the above frameworks by defining 16 AI use activities that are technology-agnostic and domain-agnostic. Any design task may consist of one or more of these activities, and each activity can be performed by a human, an AI system, or a human–AI team. Table 1 gives examples of selected activities relevant to engineering and water infrastructure planning.

Table 1: Selected human–AI use activities from the NIST taxonomy (National Institute of Standards and Technology, 2024)

Activity	Description	Example output
Image analysis	Identifying features in digital images to extract meaningful information	Medical diagnostics
Information retrieval	Finding information on specific topics of interest	Accelerating the search for stable proteins for drug development
Monitoring	Observing and tracking a process, quality, or state over time	Wildfire monitoring
Prediction	Anticipating the probability of a future outcome	Sales forecasting; weather prediction
Process automation	Executing repetitive tasks, reducing bottlenecks and errors	Automation of administrative tasks
Recommendation	Suggesting a manageable set of actionable alternatives to support decision-making	Content recommendations; procurement suggestions

It should be noted that existing taxonomies, including NIST, were primarily developed with more traditional AI systems in mind. The classification of large language models within these frameworks is challenging due to their rapid and continuous development. LLMs are therefore best analysed separately, using the contextual OECD framework, while NIST and CSET are more informative for characterising purpose-built AI tools in design workflows.

Summary

AI refers to software systems capable of data interpretation, decision-making, and human-like reasoning without fully pre-specified rules. From its origins in symbolic rule-based systems in the 1950s, AI has evolved through machine learning to deep learning and, most recently, to large-scale generative models based on the transformer architecture.

Machine learning encompasses supervised, unsupervised, reinforcement, and deep learning paradigms, each suited to different tasks. Deep learning models—built from multi-layer neural networks—underpin modern applications ranging from speech recognition to large language models such as ChatGPT and Claude.

Three complementary frameworks are useful for classifying AI systems in engineering contexts. The OECD framework offers a contextual, use-case-oriented perspective particularly suited to LLMs. The CSET autonomy taxonomy describes how independently a system acts. The NIST use taxonomy defines 16 technology-agnostic use activities applicable across design tasks, though it is better suited to traditional AI tools than to LLMs.

In water infrastructure planning, AI offers genuine opportunities in scenario generation, network optimisation, predictive analytics, and expert support. The principal barriers to adoption are currently data scarcity and fragmentation, data security constraints, a lack of purpose-built design tools, and the limited interpretability of complex models. AI is best understood as a tool to augment the engineer’s capabilities rather than to replace professional expertise.

Claude.ai was used to help translate the content of PowerPoint slides to english, as the final deliverable to Sweco was a PPT presentation in Finnish, not a report.

Frameworks

Theory of Change: Logic and Signal and Signal Based Strategy

To systematically analyze the integration of artificial intelligence within established design processes, this report employs the Theory of Change Logic Model framework (Funnell and Rogers, 2011). This framework provides a structured causal pathway for evaluating organizational transformation, dividing the adoption process into two primary phases: the foundational 'Plan' (comprising Resources and Activities) and the subsequent 'Desired Results' (comprising Outputs, Outcomes, and ultimate Impact). In the context of engineering consultancies such as Sweco, current organizational assessments indicate that the industry is still positioned at the beginning of this chronologi-

cal continuum. Consequently, the immediate focus lies heavily on identifying and overcoming initial operational bottlenecks, such as time constraints and a lack of practical reference cases among designers. Utilizing this logic model allows the report to systematically trace how mobilizing essential early-stage resources (e.g., designer competencies, project data) and executing targeted activities (e.g., pilot projects, training) are critical prerequisites for eventually realizing long-term strategic impacts, such as the automation of routine tasks, enhanced overall productivity, and increased meaningfulness of expert work.

According to Wack (1985) in his analysis of Royal Dutch/Shell's strategic planning, attempting to predict a single future through traditional economic forecasting is highly inadequate in a turbulent business environment where variables constantly shift. Instead of seeking precise predictions, organizations should construct a flexible strategic roadmap based on multiple scenarios that explore the interplay between predetermined elements and critical uncertainties (Wack, 1985). This methodology prepares managers for various plausible realities by identifying specific signposts or signals. As Wack illustrates through the analogy of a canoeist hearing white water before encountering rapids, strategic foresight relies on recognizing the early signals that prefigure major environmental disruptions. By actively monitoring the external environment for these strategic indicators, decision makers can recognize exactly which scenario is unfolding in the real world. Consequently, this approach systematically recalibrates the internal mental models, or microcosms, of executives, allowing them to abandon obsolete assumptions and adapt their strategies dynamically as actual events reveal the true trajectory of the (Wack, 1985). In Sweco's case this means monitoring signals from customers, software vendors and inside the company to see how AI is changing the engineering consulting landscape in water network planning.

Automation–Augmentation–Autonomy

The transition to the Automation–Augmentation–Autonomy framework was derived from the Boston Consulting Group (2026) analysis by shifting the focus from "jobs" to "tasks" and their specific relationship with human expertise. While traditional models often view AI as a binary tool for replacement, the BCG text emphasizes that AI's impact is a spectrum: it replaces routine, structured actions (automation), enhances human productivity in high-value roles through collaborative feedback loops (augmentation), and increasingly manages end-to-end goal-directed processes (autonomy). By synthesizing these dynamics, the study was able to move beyond a simple efficiency metric and instead create a categorical lens to evaluate where water supply designers retain their "silent knowledge" and where they can safely delegate decision-making authority to an intelligent system.

3 Data & Methods

Expert interviews

To understand the current water supply general planning process, to identify potential applications for artificial intelligence (AI), five expert interviews were conducted. The participants represented various roles within the planning lifecycle, including a project manager, general and construction designers, an data analytics expert, and an AI trainer.

Research Objectives

The primary goal of these interviews was to map the detailed stages of the water supply planning process from a designer’s perspective. Specifically, the study sought to answer:

- In which stages of the general planning process can AI provide tangible value?
- Where does the core value of human expertise remain indispensable?
- What are the primary time-consuming challenges where value is created today?
- What are the software tools that designers and project managers use most often?
- What are the boundaries of acceptability and responsibility when comparing AI roles (automation, augmentation, and autonomy) to manual tasks?

Methodological Frameworks

The analysis and interviews were structured using three distinct theoretical frameworks:

- **Planning Process and Software Tools:** This framework was used to locate where AI potential exists within current workflows and the software tools, such as CAD, GIS, and modeling software, used by designers.
- **Time Allocation and the Core of Expert Work:** This lens focused on identifying where value is generated and where professional “silent knowledge” (the ability to assess practical feasibility in the field) is most critical.
- **Automation–Augmentation–Autonomy (adapted from Boston Consulting Group, 2026):** This framework categorized AI’s poten-

tial roles based on the degree of decision-making and the division of labor between humans and machines:

- **Automation:** AI performs clearly defined, repetitive, and mechanical tasks to speed up routines and free up expert time.
- **Augmentation:** AI acts as a “sparring partner,” providing new perspectives for decision support while the human maintains full control.
- **Autonomy:** AI operates independently on entire task entities, making decisions within predefined constraints without constant human guidance.

Software Interviews

In addition to expert interviews, four representatives from three software vendors were interviewed to gain insight into their perspectives on the use of AI in water network planning and the consulting business, including how they are preparing for its adoption, what new features or software solutions are under development, and what risks they foresee. Software vendor interview uses the same framework as expert interviews.

Research Objectives

- What kind of AI features and capabilities do they have know?
- How they are integrating AI into their software tools and what type of AI is it e.g. machine learning (ML) or deep learning (DL) etc.?
- What category does it fall into, traditional software automation, augmentation or autonomy?
- How AI is going to change consulting business in water network planning?
- What risks do they foresee in implementing AI as part of the design process?

Workshop

The final method consists of a workshop designed to present the results to date, with a research objective to validate the mapped process and gather expert feedback on the scenarios. The workshop is structured into three parts: first, participants complete and refine the process map by identifying missing tasks; second, they assess which tasks can be automated, augmented, or fully autonomised now or within a 3–5 year horizon; third, they evaluate

the scenarios and provide estimates of their likelihood and expected timing, if any.

Research objectives

A workshop was organized to validate the mapped planning process and gather expert perspectives on scenarios describing how AI could impact the water network planning consulting industry.

- Identify and add missing tasks to the process map.
- Map tasks that could be automated, augmented, or autonomized either currently or within the next 3 to 5 years.
- Gather expert views on the likelihood of the proposed scenarios and their expected timeframe for realization.

4 Results

AI in Water Infrastructure Planning Based on Literature

The Nature of Water Infrastructure Design

Water infrastructure planning involves integrating objectives, resources, and uncertainty management. In practice, it encompasses decisions about network dimensioning and routing, determining pipe routes and pump station locations, and operational planning under changing demand and supply conditions (Shayboun et al., 2020; Li et al., 2025; Sweco Finland Oy, 2026).

Water network problems can be divided into *design problems* (pipe routing, pump station siting) and operational problems (managing flow under variable demand). At its most elemental, design resembles the act of drawing a pipe route on paper while observing constraints such as protected natural areas.

General planning in water infrastructure can be organised across several levels of spatial scope: regional plans at the provincial and sub-regional scale (covering water supply, emergency preparedness, and trunk sewers); municipal plans (network modelling, expansion scheduling); neighbourhood-level planning and water cooperative schemes; and component-level assessments (rehabilitation options for pipes and pump stations based on condition and material). Because water infrastructure planning spans such diverse decision types, the application of AI must also be considered from multiple perspectives rather than evaluated at the level of individual tools.

AI Opportunities in Water Infrastructure Planning

AI offers several opportunities to strengthen engineering design workflows in the water sector (Sweco Finland Oy, 2026; Taiwo et al., 2025; Gacu et al., 2025; Almulhim, 2025):

Scenario generation and demand forecasting. AI can generate a diverse range of demand scenarios to support water supply and distribution planning, facilitating long-term planning under uncertainty.

Network optimisation. AI-based optimisation can simultaneously balance network efficiency and resilience, particularly over long planning horizons.

Critical scenario review. AI can assist engineers in critically evaluating scenarios, surfacing novel perspectives and supporting expert judgement rather than replacing it.

Tool customisation. Engineers can already use AI to automate small but time-consuming tasks—for example, by using language models to write macros for existing planning software.

Predictive analytics. AI enables more proactive decision-making and resource allocation, including early warning for water quality degradation, flood risk, and network load (Gacu et al., 2025; Almulhim, 2025).

Digital twins. AI supports the creation and operation of digital twin models for real-time resource optimisation (Almulhim, 2025).

Limitations and Constraints

Despite its promise, the large-scale adoption of AI in water infrastructure design faces several constraints (Sweco Finland Oy, 2026; Byrge et al., 2025; Taiwo et al., 2025):

Data availability and quality. AI models require substantial quantities of high-quality, standardised, and accessible data. In the water sector, source data may be unavailable, fragmented across systems, or insufficiently standardised for integration.

Data security. Practitioners report that data security requirements prevent source data from being held in a single location, which complicates the aggregation needed for AI model training.

Lack of suitable tools. The primary constraint on large-scale uptake is currently the absence of purpose-built tools for water network design specifically. Existing research tends to focus on network management and project scheduling rather than on design tasks.

Interpretability. The limited interpretability of AI models—particularly

deep learning systems—can undermine trust, especially in high-stakes applications where the reasoning behind a recommendation must be transparent and auditable (Gacu et al., 2025).

Geospatial reasoning. AI systems may struggle to reason accurately about geospatial data, which is central to infrastructure routing and siting decisions.

Absence of genuine creativity. AI cannot produce genuinely new ideas grounded in independent understanding in the way human experts can (Byrge et al., 2025). It is therefore better suited to supporting and augmenting the designer’s work than to replacing human judgement.

Unrecognised opportunities. Practitioners do not yet fully recognise the range of tasks for which AI could provide meaningful support, which may slow adoption even where capable tools exist.

Expert interview results

The results of the interviews with Sweco experts highlight a clear distinction between technical potential and the practical boundaries of professional responsibility in water infrastructure planning. The findings are categorized into immediate technical opportunities and the critical role of human oversight.

Potential AI Applications in the Planning Process

The interviewees identified several specific areas where AI could improve operational efficiency, particularly in repetitive, structure-heavy tasks where errors typically stem from oversight rather than a lack of expertise:

- **Data Automation:** The most significant immediate benefit is the automated gathering and collation of decentralized source data, such as GIS, land use plans, soil data, and protected areas from various authorities.
- **Documentation and Reporting:** AI can be used to generate routine project descriptions, such as soil, waterway, and nature conservation reports, which often follow a repetitive structure across different projects.
- **Technical Tasks:** Tools can assist in creating meeting minutes from transcripts, automating macros for mass and line length calculations, and performing formal plan validation (e.g., checking fonts and layers).
- **Augmentation:** AI is viewed as a “sparring partner” that can validate or criticize alignment options, combine the best parts of different

design alternatives, and simplify technical text for stakeholder communication.

Expert Insights on Responsibility and the “Black Box” Problem

Despite the technical possibilities, the interviews revealed a strong consensus regarding the limits of AI, particularly concerning autonomy in critical infrastructure:

- **Professional Liability:** The designer bears the ultimate legal and professional responsibility for the plan. Interviewees expressed that they cannot trust AI as much as their own expertise.
- **The Verification Paradox:** A central concern is that reviewing a “black box” plan produced by AI might take more time than creating it from scratch, as the designer must verify every detail before signing off.
- **Silent Knowledge:** AI currently lacks the “silent knowledge”: the ability to evaluate from a map whether a solution is practically feasible in the field, considering factors like terrain accessibility or existing vegetation.
- **Human-Centric Decisions:** Tasks such as final alignment choices, risk assessments, and local contextual deliberation must remain under human control.

Barriers to Adoption and Organizational Change

The primary obstacles to adopting AI are not a resistance to technology itself, but rather practical constraints. These include a lack of allocated time for designers to experiment, the abstract nature of current instructions, and data security concerns related to critical infrastructure. The experts suggested that a successful transition requires a low-threshold culture of experimentation and practical examples that relate directly to their daily design context.

Software Vendor Interview Results

Interviews with software vendors indicated that plans to integrate AI into existing tools are underway, although implementation is still at an early stage. Current development efforts focus primarily on automation, with some augmentation-oriented use cases, such as automated satellite image recognition using machine learning and route-planning support demonstrated by Deutsche Telekom. Similarly to expert interview, data security, quality and

completeness were identified as the primary risks associated with AI implementation. According to the interviewees, the prevalent hour-based billing model in water network planning may need to evolve if planning processes become significantly automated.

- The primary driver for implementing new features is customer demand, highlighting the importance of proactive collaboration between Sweco and software vendors.
- Data security was the main concern with cloud-based AI tools.
- Deutsche Telekom demonstrated an ML-based algorithm for identifying a cost-optimal fiber optic cable route ((Chaos Computer Club, 2024)). One central use case identified by a software vendor was the use of machine learning to analyze satellite imagery for automatic object recognition, such as detecting buildings and mapping them spatially.
- Data quality and completeness were considered critical, as insufficient or incomplete data significantly reduces output quality.

Planning process in general water network

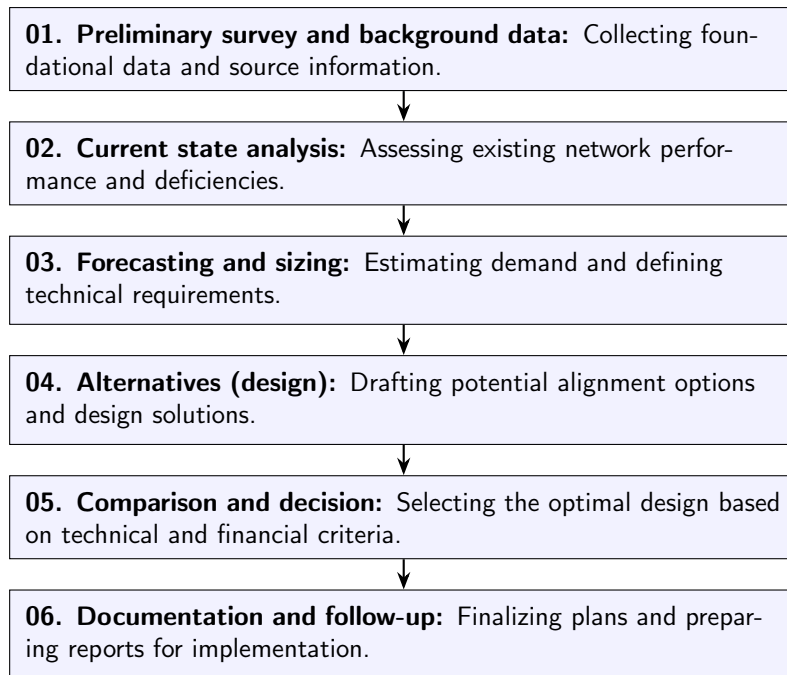


Figure 1: The six-phase water network planning process.

The general water network planning process is a six-phase workflow that transforms raw site information into a final engineering design. The chart in Figure 1 illustrates the process.

Mock-up: AI-Assisted Source Data Synthesis

This section outlines a conceptual mock-up designed to automate the collection and analysis of spatial data in infrastructure planning. The primary objective is to replace the current manual process (which often requires several hours of gathering data from various portals (e.g., SYKE, Finnish Heritage Agency) and compiling it in GIS software) with an integrated AI-driven solution. An mock-up of the tool was created to help Sweco visualize the tool.

Functional Overview and Technical Implementation

The tool utilizes an interface where a designer defines a study area on a map. The system then executes automated API calls to retrieve data via two primary interface types:

- **WFS (Web Feature Service):** This service returns vector data (e.g., groundwater polygons, conservation points) in formats such as GeoJSON or GML. This structured data allows the system to perform geometric analysis, such as calculating surface areas or identifying intersections with the planning area, which are subsequently converted into text for a Retrieval-Augmented Generation (RAG) index.
- **WMS (Web Map Service):** This service provides pixel-based images. Although suitable for visualization, this format requires vision-based AI models to interpret imagery into processable text for deeper analysis.

The final output provides an interactive map visualization, an automated textual summary (e.g. quantifying groundwater areas or identifying archeological sites), and a Q&A assistant to help designers query specific project details.

Data Integration and Sources

The tool integrates diverse data streams categorized by availability and source:

1. **Public Open Data:** Includes soil, groundwater, and nature maps from Paikkatietoikkuna and SYKE, cultural heritage data from the Finnish Heritage Agency, and population projections from Statistics Finland.
2. **Project-Specific Data:** Incorporates client-provided information (e.g., municipal utility maps) and internal reference archives of previous designs.

Critical Risks and Technical Uncertainties

The implementation of AI in critical infrastructure planning involves several technical challenges.

- **Data Integrity:** Risks include "blind spots" (missing or classified data), varying update frequencies between different agencies, and potential coordinate system discrepancies when merging local datasets with national standards.
- **AI Reliability:** There is a risk of RAG-related "hallucinations," where the AI may synthesize unrelated information into a misleadingly plausible summary under time pressure.
- **Security:** Sensitive data, particularly regarding water supply and utility networks, requires a rigorous evaluation of cloud-based versus on-premise deployment to ensure information security.

Scenarios

Four distinct scenarios were created to map out the potential futures for water network planning consulting industry based on AI usage and degree of work automation, ranging from business-as-usual with no major changes to total disruption where customers can order complete plans and almost work is automated and humans focus on validation.

- **Scenario 1: Business as usual.** Limited, ad hoc task automation without systematic deployment. Signals include rising productivity, downward pricing pressure, and client expectations of faster delivery.
- **Scenario 2: Software vendor AI rollout.** Established software vendors integrate AI to automate and augment design tasks. This increases productivity and intensifies pricing pressure. Signals include M&A activity between vendors and data providers, increased media coverage of AI in water network planning, and accelerating delivery expectations.
- **Scenario 3: Third-party AI platform emergence.** Independent AI developers provide API-based tools for design tasks, enabling in-house and hybrid solution development by consultancies. Signals include recruitment of AI talent by consulting firms, internal tool development initiatives, and growing regulatory and client acceptance of AI-generated designs.
- **Scenario 4: Full disruption by AI platforms.** Third-party AI platforms enable end-to-end automated planning, significantly reducing the role of traditional consultants. Signals include declining project win rates, migration of engineers from consulting firms to AI tool

providers for development and validation roles, and regulatory frameworks increasingly supporting AI-based design outputs.

Workshop Results

The process was largely mapped, with only a few gaps remaining, such as permitting tasks handled by the Centres for Economic Development, Transport and the Environment (ELY-keskus). Tasks identified as suitable for automation included data collection from public sources. Augmentation opportunities included route generation and selection of optimal alternatives. Autonomous execution was considered feasible for document updates in response to plan changes.

When asked to estimate the timeframe of each scenario, respondents viewed Scenarios 1 and 2 as current reality. Scenario 3 was considered likely within 1–3 years. Scenario 4 was estimated at 3–5 years under favourable regulatory conditions, or 10 years or longer, potentially never, if regulation remains restrictive.

5 Discussion

This study explored the integration of Artificial Intelligence into Sweco's water network planning processes, revealing a significant, yet largely untapped, potential to enhance efficiency. The findings indicate that while current design workflows are heavily expert-driven and manual, many routine tasks such as data collection and initial report generation have most potential for automation. The expert interviews confirmed that the most immediate value of AI lies in automating these time-consuming, repetitive tasks, thereby freeing up engineers to focus on higher-value activities like complex problem-solving and stakeholder engagement.

The developed scenarios highlight a critical strategic crossroads for Sweco and the broader consulting industry. The "Business as usual" and "Software vendor AI rollout" scenarios (Scenarios 1 and 2) represent the current trajectory, where productivity gains are incremental and largely driven by external software providers. However, Scenarios 3 and 4, which involve the emergence of third-party platforms and potential market disruption, underscore the need for proactive strategy. For Sweco, this signals a potential shift in its business model, moving from hourly billing to delivering projects with fixed pricing. This transition requires not only new technological capabilities but also a cultural shift towards embracing AI, starting with small pilot groups and expanding. Monitoring signals, such as the recruitment of AI talent by competitors, designers leaving Sweco or the launch of new API-based design tools and regulation around AI use in water network planning, will be crucial for navigating this evolving landscape.

Limitations include the project team’s limited subject-matter expertise, requiring reliance on expert interviews rather than quantitative data, which may introduce bias. The sample size is small (five expert interviews and three software vendor interviews), and role coverage is narrow, limited to designers, project managers, and an AI trainer, which restricts representativeness. The scenarios are exploratory rather than predictive and are not based on data indicating likelihood. They should be seen as structured future possibilities used to frame discussion, not as forecasts or probability-based outcomes.

6 Conclusions

In conclusion, this project found that Sweco is in the early stages of adopting AI in water network planning, with substantial potential for improving operational efficiency. The investigation revealed that the greatest opportunities for immediate impact are in the automation of data synthesis and the augmentation of the design process, rather than the autonomous execution of entire planning tasks. Human expertise remains indispensable, particularly for validating results, navigating complex local constraints, and assuming professional liability.

The path forward for Sweco involves a strategic, phased approach to AI integration. The roadmap developed through this project recommends beginning with low-risk, high-impact applications. A key recommendation is to pilot an AI-assisted tool for automating the collection and synthesis of decentralized source data, which was identified as a major bottleneck in the current workflow.

Regardless of which future scenario materializes, certain actions are universally beneficial. Sweco must invest in building a solid data infrastructure and enhancing the AI literacy of its staff. Fostering a culture that encourages experimentation with new tools is essential for organic capability growth. Proactive engagement with software vendors and clients is also critical to steer the development of new tools that meet the specific needs of water infrastructure planning.

Sweco already possesses its most valuable resources: deep domain expertise and extensive historical project data. To leverage these assets in the age of AI, the organization will need to supplement them with dedicated resources for innovation, strategic partnerships, and a clear vision for integrating AI as a core component of its service delivery. Ultimately, the successful adoption of AI will enable Sweco to transition from a provider of engineering services to a purveyor of data-driven insights and optimized solutions.

References

- Catherine Aiken. Classifying AI systems, 2021. URL <https://cset.georgetown.edu/publication/classifying-ai-systems/>. Accessed: March 2026.
- Abdulrahman I. Almulhim. Integrating artificial intelligence into smart infrastructure management for sustainable urban planning. *Technologies*, 13(11):481, 2025. doi: <https://doi.org/10.3390/technologies13110481>.
- Paolo Bory, Simone Natale, and Christian Katzenbach. Strong and weak AI narratives: An analytical framework. *AI & Society*, 40(4):2107–2117, 2024. doi: <https://doi.org/10.1007/s00146-024-02087-8>.
- Boston Consulting Group. AI will reshape more jobs than it replaces, 2026. URL <https://www.bcg.com/publications/2026/ai-will-reshape-more-jobs-than-it-replaces>. Accessed: May 2026.
- Christian Byrge, Erik E. Guzik, and Christian Gilde. Artificial intelligence and the creative process: Does AI-creativity extend beyond divergent thinking? *Journal of Creativity*, 35(2):100105, 2025. ISSN 2713-3745. doi: <https://doi.org/10.1016/j.yjoc.2025.100105>.
- Chaos Computer Club. How digging in the earth for the fibre roll-out took grass to the cloud. <https://media.ccc.de/v/bucharest-234-how-digging-in-the-earth-for-the-fibre-roll-out-took-grass-to-the-cl>. 2024. Accessed: 2026-05-28.
- Blagoj Delipetrev, Chrysi Tsinaraki, and Uros Kostic. Historical evolution of artificial intelligence, 2020. EUR 30221 EN, JRC120469. doi: <https://doi.org/10.2760/801580>. Accessed: March 2026.
- Jerome G. Gacu, Cris Edward F. Monjardin, Ronald Gabriel T. Mangulabnan, Gerald Christian E. Pugat, and Jerosse G. Solmerin Solmerin. Artificial intelligence (AI) in surface water management: A comprehensive review of methods, applications, and challenges. *Water*, 17(11):1707, 2025. doi: <https://doi.org/10.3390/w17111707>.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, Cambridge, MA, 2016.
- Ran He, Jie Cao, and Tieniu Tan. Generative artificial intelligence: A historical perspective. *National Science Review*, 12(5):nwaf050, 2025. ISSN 2095-5138. doi: <https://doi.org/10.1093/nsr/nwaf050>.
- IBM. What is deep learning?, 2024a. URL <https://www.ibm.com/think/topics/deep-learning>. Accessed: March 2026.
- IBM. What are large language models?, 2024b. URL <https://www.ibm.com/topics/large-language-models>. Accessed: March 2026.

- Rich Kenny, Avinash Lunj, and Alexandra Kis. AI’s thirst for water, 2025. URL <https://sustainableict.blog.gov.uk/2025/09/17/ais-thirst-for-water/>. Accessed: March 2026.
- Yijiang Li, Santanu S. Dey, Nikolaos V. Sahinidis, Naresh Susarla, Miguel A. Zamarripa, and Markus G. Drouven. Optimizing the design and operation of water networks: Two decomposition approaches. *Computers & Chemical Engineering*, 192:108897, 2025. ISSN 0098-1354. doi: <https://doi.org/10.1016/j.compchemeng.2024.108897>.
- MLOps Community. The impact of prompt bloat on llm output quality, 2026. URL <https://mlops.community/the-impact-of-prompt-bloat-on-llm-output-quality/>. Accessed: March 2026.
- National Institute of Standards and Technology. AI use taxonomy, 2024. NIST.AI.200-1. U.S. Department of Commerce, Gaithersburg, MD. doi: <https://doi.org/10.6028/NIST.AI.200-1>. Accessed: March 2026.
- NVIDIA. What are large language models?, 2026. URL <https://resources.nvidia.com/en-us-new-llm-product/large-language-models-overview-web-page>. Accessed: March 2026.
- OECD. OECD framework for the classification of AI systems, 2022. URL https://www.oecd.org/en/publications/oecd-framework-for-the-classification-of-ai-systems_cb6d9eca-en.html. Accessed: March 2026.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986. doi: <https://doi.org/10.1038/323533a0>.
- May Shayboun, Dimosthenis Kifokeris, and Christian Koch. Construction planning with machine learning. In *Proceedings of the 35th Annual ARCOM Conference*, pages 699–708, Leeds, United Kingdom, 2020. Association of Researchers in Construction Management (ARCOM). ISBN 978-0-9955463-4-9.
- Sweco Finland Oy. Internal material and expert interviews on AI in water infrastructure design. Unpublished internal document, 2026.
- Ridwan Taiwo, Abdul-Mugis Yussif, and Tarek Zayed. Making waves: Generative artificial intelligence in water distribution networks: Opportunities and challenges. *Water Research X*, 28:100316, 2025. ISSN 2589-9147. doi: <https://doi.org/10.1016/j.wroa.2025.100316>.
- Samuel Tatipamula. The evolution of AI workflow automation: From rules to

adaptive intelligence. *International Journal of Science and Advanced Technology*, 16(2), 2025. doi: <https://doi.org/10.71097/IJSAT.v16.i2.3531>.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, 2017.

7 Self Assessment

Following the Plan

We initially followed the original project plan closely. However, delays in completing the first two work packages caused subsequent work packages to also fall behind schedule, leaving less time for later stages of the project. The project scope remained largely consistent with the original description, although software vendor interviews and scenario analysis were added to broaden the perspective of the study. Overall, the project description provided a clear framework from the beginning, allowing the work to focus primarily on execution. Finally, project was successful because we worked on it steadily for the entire time making atleast some progress weekly.

Project Success

Overall, the customer was very happy with the workshop and the results such as the roadmap and the report, because it elicited interest and was very thought provoking. Moreover, the customer ended up liking scenario-based thinking, despite being a backup plan. Literature review was sufficiently broad and detailed to cover the background. Choosing appropriate frameworks, particularly the automation–augmentation–autonomy classification, enabled a more systematic approach to the subject matter and improved the validity of the study.

Failures

The project faced several challenges during execution. The original schedule was not maintained, and as a result, the fourth work package was likely not as comprehensive as it could have been. In retrospect, the project could also have included more concrete examples of AI applications instead of focusing mainly on conceptual ideas for where AI could be utilized within the planning process. Additionally, the overall coherence between different sections of the project became somewhat unclear over time. Communication within the team was at times ineffective, and shared deadlines were not always met consistently. There were also periods during the project when the next steps were unclear, despite there being available time to advance the

work further. Furthermore, future objectives could have been broken down into more actionable tasks and allocated more clearly among team members by the project manager in order to reduce idle time and improve workflow continuity. In addition, the project could have benefited from stronger involvement from the customer organization through more active feedback and idea generation, rather than relying primarily on guidance from the Aalto supervisor.

Improvements

Several lessons were also identified regarding teamwork and project management. Tasks were often left until the last moment, and iteration cycles within the team were relatively slow. Although the work was ultimately completed successfully, the process could have been more efficient. Communication within the team was repeatedly discussed as an area for improvement, but these discussions did not always translate into practical changes. At the beginning of the project, roles and responsibilities could also have been defined more clearly, for example by assigning each member primary responsibility for a specific work package. Instead, responsibilities were shared broadly, which on one hand allowed everyone to gain a wider understanding of the project, but on the other hand may have reduced ownership and initiative in driving individual sections forward.

In retrospect, team members could also have taken more independent responsibility earlier in the project rather than relying heavily on feedback from the designated responsible persons such as. While teamwork naturally includes independent tasks, an effective team environment also requires active engagement with others' work through feedback and collaborative improvement. At times, it felt as though some team members were primarily focused on completing only their own assigned contributions.

Customer could have been more involved by providing more feedback so that the project could have been more catered to their needs.