Mälardalen University Licentiate Thesis 381

Towards a Digital Twin For Quarry Sites:

From Requirements to Operational Components

Abdulkarim Habbab

Mälardalen
University

Mälardalen University Press Licentiate Theses
No. 381

TOWARDS A DIGITAL TWIN FOR QUARRY SITES: FROM REQUIREMENTS TO OPERATIONAL COMPONENTS

Abdulkarim Habbab

2026

Department of Computer Science & Engineering

ISBN 978-91-7485-749-8
ISSN 1651-9256
Printed by E-Print AB, Stockholm, Sweden

Abstract

Off-road quarry environments are complex systems where machines, materials, and humans interact under harsh, safety-critical, and resource-constrained conditions. While digital twins promise to transform these operations through virtual experimentation and optimization, practical adoption is limited by three key challenges: integrating models across multiple temporal and spatial scales, balancing computational efficiency with physical fidelity, and maintaining modular architectures that can evolve with changing site requirements.

This licentiate thesis contributes toward addressing these challenges by developing a foundational framework for digital twin implementation in quarry operations and demonstrating its feasibility through two enabling components. First, through industry-embedded case studies combining semi-structured interviews, expert workshops, and site observations, the research maps simulation-optimization requirements across three operational levels and proposes a hierarchical modeling framework that defines interfaces between site-level planning, operational coordination, and machine dynamics. This framework establishes how information should flow between high-level production scheduling and low-level equipment dynamic and control while maintaining computational tractability.

To demonstrate technical feasibility within this framework, two machine-learning components at the dynamics level were developed. A torque-prediction model uses expert-guided feature selection and Shapley Additive exPlanations (SHAP) analysis to achieve high-fidelity estimates with minimal sensor inputs, providing a template for interpretable surrogate modeling. A Long Short-Term Memory (LSTM) based world model enables efficient reinforcement learning for autonomous bucket filling, showing major improvements in both productivity and energy efficiency compared to baseline controllers in simulation environments.

This research establishes the architectural foundation and demonstrates core technical capabilities necessary for quarry digital twins, while explicitly deferring full system integration, field validation, and cross-site deployment to future doctoral work. The contributions provide a structured approach to multi-level modeling for quarry digital twins, establishing methodological foundations for integrating site level planning, operational coordination, and machine dynamics models while demonstrating that machine learning can deliver computationally efficient surrogates suitable for real-time applications.

Sammanfattning

Terrängbaserade täktmiljöer är komplexa system där maskiner, material och människor samverkar under krävande, säkerhetskritiska och resursbegränsade förhållanden. Även om digitala tvillingar har potential att omvandla dessa verksamheter genom virtuella experiment och optimering, begränsas den praktiska användningen av tre centrala utmaningar: att integrera modeller över flera tids- och rumsskalor, att balansera beräkningseffektivitet mot fysisk noggrannhet samt att upprätthålla modulära arkitekturer som kan utvecklas i takt med förändrade krav på anläggningen.

Denna licentiatuppsats bidrar till att hantera dessa utmaningar genom att utveckla ett grundläggande ramverk för implementering av digitala tvillingar i täktverksamhet och genom att demonstrera dess genomförbarhet med hjälp av två möjliggörande komponenter. För det första kartlägger forskningen, genom industrinära fallstudier som kombinerar semistrukturerade intervjuer, expertworkshoppar och platsobservationer, krav på simulering och optimering över tre operativa nivåer och föreslår ett hierarkiskt modelleringsramverk som definierar gränssnitt mellan planering på anläggningsnivå, operativ samordning och maskindynamik. Ramverket fastställer hur information bör flöda mellan övergripande produktionsplanering och låg nivåns styrning av utrustning, samtidigt som beräkningsmässig hanterbarhet bibehålls.

För att demonstrera teknisk genomförbarhet inom detta ramverk utvecklar uppsatsen två maskininlärningskomponenter på dynamiknivån. En momentprediktionsmodell använder expertstyrt urval av egenskaper (features) och SHapley Additive exPlanations (SHAP)-analys för att uppnå högupplösta skattningar med minimalt antal sensorindata, vilket ger en mall för tolkbara surrogatmodeller. En världsmodell baserad på Long Short-Term Memory (LSTM) möjliggör effektiv förstärkningsinlärning för autonom skopfyllning och visar stora förbättringar i både produktivitet och energieffektivitet jämfört med baslinjestyrningar i simuleringsmiljöer.

Denna forskning etablerar den arkitektoniska grunden och demonstrerar centrala tekniska förmågor som krävs för digitala tvillingar i täkter, samtidigt som fullständig systemintegration, fältvalidering och driftsättning över flera anläggningar uttryckligen lämnas till framtida doktorandarbete. Bidragen ger ett strukturerat angreppssätt för flernivåmodellering av digitala tvillingar för täktverksamhet, etablerar metodologiska grunder för att integrera planering på anläggningsnivå, operativ samordning och modeller för maskindynamik, samt visar att maskininlärning kan leverera beräkningseffektiva surrogat som är lämpliga för realtidsapplikationer.

To my Family

Acknowledgments

Reaching the halfway point of my PhD with this licentiate thesis feels like looking back on a journey full of growth, challenge, and discovery. I've been incredibly lucky to share this path with many wonderful people who have made it not only possible but deeply rewarding.

My warmest thanks go to my supervisors, Markus Bohlin, Anas Fattouh, and Koteshwar Chirumalla, whose trust and guidance gave me both the freedom and direction I needed to grow as a researcher.

I would like to express my sincere appreciation to Volvo Construction Equipment for the support that enabled this industrial collaboration and made the research in this thesis possible. I am particularly grateful to my industrial supervisor, Bobbie Frank, and my mentor, Elianne Lindmark, for their continuous guidance, insightful feedback, and constructive collaboration throughout the project. I also wish to thank Anna Sannö and Per Verner for opening the door to this research area—Anna for helping to establish the conditions for this work, and Per for welcoming me into his team as I embarked on the PhD journey.

To my fellow PhD colleagues at Volvo CE (Phillip, Nicolas, Barret, Natalie) and at MDU (Sarmad, Abbas, Edin, Enxhi): thank you for the discussions, help, and laughter that made even the toughest weeks more enjoyable.

This work was partially funded by the KKS INDTECH Industrial School (Grant No. 20200132 01 H), and by Volvo Construction Equipment AB through the projects TRUST-SOS (TRUSTed Site Optimization Solutions; Ref. 2021-02551; Nov 2021–Dec 2024) and TESTED-SOS (Tested Site Optimization Solutions; Ref. 2024-03678; Nov 2024–Sep 2027).

Abdulkarim Habbab, Eskilstuna, March 2026

List of Publications

Papers included in this thesis[1]

Paper A: *Abdulkarim Habbab*, Anas Fattouh, Bobbie Frank, Koteshwar Chirumalla, Markus Bohlin. "*Mapping simulation optimization requirements for construction sites: A study in the heavy-duty vehicles industry*". In the 64th International Conference of Scandinavian Simulation Society (SIMS 2023), Västerås, Sweden.

Paper B: *Abdulkarim Habbab*, Anas Fattouh, Bobbie Frank, Elianne Lindmark, Koteshwar Chirumalla, Markus Bohlin. "*A Multilevel Modelling Framework for Quarry Site Operations*". In Proceedings of the 12th ACM/IEEE International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems (SESoS 2024), Lisbon, Portugal.

Paper C: *Abdulkarim Habbab*, Anas Fattouh, Bobbie Frank, Mohammad Loni, Koteshwar Chirumalla, Markus Bohlin. "*Efficient Torque Prediction for Digital Twins in Quarry Operations: A Data-Driven and Expert-Guided Approach*". In the IEEE International Conference on Industrial Informatics (INDIN 2025).

Paper D: Duarte Morais, *Abdulkarim Habbab*, Shehr Bano Fatima, Alexandre Proutiere. "*Reinforcement Learning with World Models for Autonomous Excavation Optimization in Wheel Loaders*". In the 66th International Conference of Scandinavian Simulation Society (SIMS 2025), Stavanger, Norway.

---

[1] The included publications have been reformatted to comply with the thesis layout.

Contents

I

Thesis

Chapter 1

Introduction

Quarry operations supply the rock, sand, and aggregates used in roads, buildings, dams, and other infrastructure projects. Operations take place under demanding conditions such as variable terrain, harsh weather, and complex logistical requirements that challenge both machines and operators daily. These operations are made up of a chain of activities like controlled blasting, excavation, loading, hauling, crushing, screening, and stockpiling, executed by fleets of heavy machines and human operators.

Today, the industry faces rising challenges that require transformation in several domains. Climate neutrality commitments mean reducing energy consumption and emissions across all operations, while competitive pressures demand improved machine effectiveness and productivity. These challenges require a better understanding of machine performance in real-world conditions. Daily decisions about fleet composition, traffic flow, and material handling affect throughput, cost, energy use, and safety across the site. Although site management teams currently maintain reasonably effective performance levels in these areas, achieving, maintaining and improving this level of operational optimality requires considerable investment in time and cognitive effort.

To meet these demands without increasing that burden, the industry is shifting toward data-driven digital services that support productivity, safety, and operational efficiency through decision support and continuous improvement [60]. Advances in digitization have changed the way quarry operations are observed and measured. Connected machines and IoT systems now stream data on position, energy, utilization, loads, and duty cycles [41, 48], while site surveys and monitoring capture terrain and stockpiles [43]. Sensor fusion, and improved connectivity have enabled real-time visibility across operations. Combined, these streams describe the site across multiple levels; from individual machine performance to workflow processes and longer-term planning decisions [45].

The increasing availability of operational data enables a transition from monitoring to integrated system representations that support analysis and optimization. This development is part of a wider Industry 4.0 transformation across heavy industries, based on the Internet of Things (IoT), artificial intelligence (AI), and advanced data analytics [29]. At the core of this transformation is the digital twin concept (DT). Originally defined within the manufacturing domain, a DT is a high-fidelity virtual replica of a physical system that is kept synchronized with its real-world counterpart throughout its life cycle, allowing monitoring, analysis, and control [53, 56]. In manufacturing context, digital twin concept supports lifecycle management, production planning, and quality control. Extending this concept to complex quarry operations, where high-level logistical events are coupled with continuous machine dynamics, positions digital twins as a natural foundation for system-wide optimization and for closing the gap between planning and operational execution [24, 45].

Early implementations (sometimes referred to as digital shadows [50]) have been tested for equipment monitoring and site coordination, but fully integrated systems remain under development. A comprehensive digital twin would synchronize terrain models, stockpile data, traffic networks, crusher settings, and machine behavior in near real time [43, 46]. The effectiveness of such a twin depends on the fidelity and adaptability of its models: they must reflect evolving site realities, including terrain, material properties, and machine–environment interactions. Therefore, recent frameworks emphasize modular architectures that enable the monitoring and integration of diverse processes and equipment at different levels [53], with demonstrated applications in teleoperation and material flow tracking [52].

Following ISO 23247 [2], which defines a digital twin as a fit-for-purpose digital representation synchronized with observable manufacturing elements, this thesis adopts a domain-specific definition tailored to quarry sites. In the quarry context, the twin is used to represent site operating conditions, including machine dynamics and machine–environment interactions, so that site-level performance can be analyzed and improved under realistic conditions. The same capability is also used to support product development, by enabling simulation and comparison of products, components, and control strategies across a range of scenarios for validation and performance assessment.

While ISO 23247 provides a general reference framework for digital-twin interoperability for manufacturing systems, the framework proposed in this thesis focuses on the modeling structure of quarry sites, its operations and machines, where processes are continuous, equipment interactions are complex, and digital twins are used to coordinate operations, manage site flow, and support domain-specific optimization efforts. It is organized into three levels of modeling which are: site, operations and dynamics, reflecting the different time scales and fidelity degree required in quarry production. Each level provides a model view and exchanges information with the others (states, events, constraints, and KPIs) to support monitoring, analysis, prediction, and decision support. In ISO 23247 terms [2], these models constitute the digital representation of the quarry system and remain meaningful through explicit data links and update mechanisms to the observed operation.

This structure aligns with the quarry's system-of-systems nature: subsystems can be engineered and operated with partial independence yet still coordinate to achieve overall production goals. Because these subsystems operate on different time scales and require models with varying abstraction and fidelity, the framework emphasizes modularity and interoperability. Models can therefore be developed, tested, and updated independently while remaining connected through defined interfaces, rather than being constrained within a single monolithic model.

Volvo Construction Equipment (Volvo CE) plays a vital role in this transformation, combining connected products and digital services to help customers act on site data [58]. Volvo CE's vision is to build integrated systems where connectivity, analytics, and equipment knowledge work together to provide decision support for quarry operations. This includes creating shared site representations for different types of analysis, synchronizing digital models with real data, and providing insights to planners and operators.

This thesis comprises a collection of studies that address digital twinning and machine learning applications in heavy equipment operations. The first study establishes a foundational requirement analysis, identifying the variables, constraints, and performance indicators necessary for site planning, operations, and machine dynamics modeling. Subsequent work builds on these requirements to propose a multi-level digital twin architecture organized into three complementary levels: site level for general performance indicators, operational level for operations description and material flow, and dynamics level for machines, materials, and machine–material interactions. This modular framework enables integration across studies while maintaining manageable computational demands.

Two distinct machine-learning components contributions are presented. One study develops a torque-prediction surrogate model combining expert-informed feature selection with SHAP analysis, demonstrating high-fidelity torque estimation using only a small number of sensor inputs. The second study investigates applying LSTM-based world models trained on simulation data to develop excavation policies for wheel loaders, achieving superior performance relative to conventional controllers in simulated environments across both productivity and energy efficiency metrics.

Together, a structured requirements map, a multi-level framework, and machine-learning components advance progress toward operational multi-scale digital twins for quarry sites. The remainder of the thesis is structured as follows. Chapter 2 covers background on quarry operations, digital twin frameworks, modeling, and optimization for quarry digital twins. Chapter 3 explains the methodology. Chapter 4 discusses the theoretical and practical implications. Chapter 5 presents the conclusions and future work. Chapter 6 presents the publications included in this thesis.

1.1 Research Gaps, Questions, and Contributions

As can be concluded from the introduction, quarry environments are complex and highly coupled systems in which high-level logistic events continuously interact with low-level physical dynamics. Although digital twin (DT) concepts promise significant transformation through virtual experimentation and optimization, the literature review in Chapter 2 highlights that existing frameworks generally fail to address this inherent "system-of-systems" nature. Site planning and machine control are often treated in isolation, limiting the potential for system-wide optimization and deep insights about the machine performance.

This thesis identifies two primary barriers that currently restrict the practical development and effectiveness of quarry digital twins:

Gap 1– **Multi-level Integration:** Current digital-twin frameworks fall short in integrating models across the temporal and spatial scales present in quarry operations, from machine-level dynamics up to site-level planning. Existing approaches either concentrate on individual machine performance or attempt site-level modeling without sufficient granularity to capture critical equipment–material or terrain interactions [32, 45, 53].

Gap 2– **Computational Efficiency vs. Model Fidelity:** Achieving high responsiveness across nonlinear dynamics like bucket-soil interaction, suspension dynamics, or crusher feed rates, while preserving physical accuracy remains challenging [25, 51]. Current simulation approaches either sacrifice fidelity for speed or require computational resources incompatible with real-time decision support.

To address these gaps, this research is guided by the following two research questions:

RQ1: How can a digital twin framework be designed to integrate quarry data and models for interpretable decision support?

RQ2: How can data-driven approaches be used to model machine-level dynamics within quarry digital twins while balancing computational efficiency, predictive fidelity, and interpretability?

The relationship between these gaps, research questions (RQs), and the contributions of this thesis (Papers A–D) is summarized in Table 1.1.

Research Gap	Research Question (RQ)	Thesis Contribution
Gap 1 (Multi-level Integration)	RQ1	Paper A, Paper B
Gap 2 (Computational Efficiency)	RQ2	Paper C, Paper D

Table 1.1: Mapping of research gaps, questions, and thesis contributions.

Chapter 2

Background

This section presents the background literature on quarry operations and digital twin technologies, outlining key industry challenges and research motivations that support the study.

2.1 The Quarry Site as an Operational Problem

Quarry operations constitute a coupled production system in which drilling and blasting, excavating, loading, hauling, crushing, screening, and stockpiling is synchronized under safety considerations, environmental, and regulatory constraints [20]. Quarry sites naturally combined of hybrid processes: discrete-events (e.g., queueing at crushers, dispatching of trucks, scheduled blasting windows) that interact with continuous machine–material dynamics (traction, bucket–pile interaction, drivetrain transients)[57]. Decisions span multiple horizons—from strategic layout and fleet sizing to shift planning and short-interval dispatch decision, while performance is assessed through multiobjective KPIs including throughput, cost, energy and fuel consumption, emissions, and safety considerations [20, 21, 57]. Early operational studies established how inefficiencies at any stage (e.g., crusher under-utilization, loader idling, haul-road congestion) propagate through the entire system, creating bottlenecks and downtime costs that influence total productivity [47, 57].

Operationally, quarries exhibit strong couplings between material flow and the utilization of resources such as machines, operators, and facilities. The balance of crusher feed is governed by the crusher's operating point and production targets, which in turn determine how loading and hauling equipment are assigned to supply it. Queuing at loading and dumping points feeds back into effective travel speeds and route choices for haul trucks. Maintenance windows and blasting schedules cause available capacity and accessible areas to vary over time. Additionally, safety rules (e.g. separation distances, speed limits, right-of-way) restrict how different types of vehicles and personnel can move and interact in shared traffic areas [19, 20]. Layout and traffic management interact with these factors: road curvature and gradients, intersection design, and stockpile placement influence achievable cycle times and energy use, while zoning and environmental permits define the allowed operating range for the site [47, 61]. Together, these interactions create a difficult planning and control problem with uncertainty and multiple objectives, where decisions that appear optimal locally (e.g., maximizing loader bucket fills) lead to suboptimal site performance. Viewed from a systems perspective, a quarry behaves as a directed system-of-systems in which multiple independently managed machines, control systems, and organizational processes must be optimized jointly rather than in isolation, further complicating site-wide optimization [8].

Human factors remain crucial to quarry operations. The skill and experience of operators directly shape how efficiently and safely they fill buckets, manage machine controls, and remain aware of changing environments [13]. Differences in operator abilities lead to variations in work speed and energy consumption. As autonomous systems become more common—whether as advanced assistance or fully automated machines—quarries will face new questions about managing mixed fleets of human-driven and automated equipment, and about how to support supervision with remote monitoring and system-level data integration [18]. This shift requires better ways to monitor operations in real time, manage risks, and ensure reliable performance [8]. From a business perspective, quarries must balance production targets with regulatory limits on emissions, noise, and dust. This creates competing priorities that shift with seasons, work schedules, and market conditions [43].

Modern quarry sites generate diverse data from machine components, sensors, surveys, and operational records. These data sources vary in sampling rates, may experience sensor drift, and have incomplete spatial coverage. Integrating this information into comprehensive representations of site conditions is essential for accurate modeling and data-driven decision-making, particularly when building computational models that should reflect real operational dynamics and context [43].

Quarry operations are complex and adaptive systems with many uncertainties and mixed human and machine elements. optimal operations depend on coordination, contextual understanding of the realities of these operations, and balancing competing goals like productivity, energy use, emissions, and safety. Understanding these systems reveals why integrated approaches are needed: deep understanding of processes and site conditions, frameworks that connect diverse data sources with operational decisions, and computational methods that balance competing priorities while adapting to real-world variability and constraints [20, 45, 61].

2.2 Digital Twin Frameworks

Supported by advances in IoT and AI, digital twin concepts (DT) are increasingly recognized as key enablers for optimizing site operations in off-road environments such as quarries, mining, and construction sites [9, 21]. These domains share operational characteristics that make them well suited for DT applications: they rely on fleets of heavy machinery (e.g. wheel loaders, haul trucks, excavators) executing material-handling workflows (drilling/blasting, excavation/loading, hauling, crushing/processing) across large, unstructured terrain, under variable weather conditions as well as safety and regulatory constraints. They also face common challenges—balancing productivity against energy costs, managing equipment wear, and adapting to dynamic site conditions (e.g., terrain evolution and material variability)—which together motivate the need for multi-level modeling and robust decision support.

This need for multi-level representations places structural requirements on DT solutions that go beyond data collection and monitoring. Therefore, recent efforts in standardization and reference-architecture emphasize the description of DT systems through explicit viewpoints and well-defined information exchanges to support consistency and interoperability [2]. Such an architecture-oriented perspective is particularly relevant in complex, data-rich operational settings, where heterogeneous models and data streams must be combined across different temporal and spatial scales.

At the conceptual level, general DT frameworks provide useful guidance for defining scope and structure. The frameworks proposed by Tao et al. [56] and Shao and Helu [53], for example, highlight the role of interfaces, synchronization mechanisms, and life-cycle considerations in aligning physical and digital systems. While valuable, these frameworks are primarily developed with controlled manufacturing environments in mind. As a result, they do not fully capture the non-deterministic physical interactions typical of quarry operations, where material variability and terrain–machine interactions introduce significant uncertainty and continuous change.

A broader synthesis of DT research is provided by Rasheed et al. [45], who survey DT value propositions, challenges, and enabling technologies across domains. In addition to multi-scale integration and the trade-off between model fidelity and computational efficiency, they emphasize cross-cutting concerns such as scalability, interoperability, security, explainability, and human–machine interaction. From a modeling standpoint, they stress the importance of real-time bidirectional data–model coupling, continual model evolution, and physically consistent, interpretable models—especially in safety-critical settings.

Against this backdrop, DT applications in heavy, off-road industries have been explored across different levels of abstraction, but multi-scale integration remains a persistent challenge. Many contributions focus either on monitoring and visualization services or on detailed dynamic studies, while system-level integration is often acknowledged, but left for future work. Consequently, modeling the entire quarry environment as a “system of systems”—linking site-level logistics and planning with operational workflows and machine dynamics—remains difficult, and most implementations stay confined to isolated operational levels [44, 56].

Examples from adjacent and related work illustrate both the promise of DTs and the limits of current practice. In civil engineering, Pregnolato et al. [43] position DTs within the “Civil Engineering 4.0” paradigm to enhance infrastructure asset management (e.g., roads and bridges). This perspective, however, typically underrepresents the short-cycle operational processes and tightly coupled machine interactions that characterize mobile quarry equipment. In the mining domain, Bertoni et al. [9] show how DTs of operational scenarios can support the design of customized Product–Service System (PSS) solutions, demonstrating strategic value but placing less emphasis on the technical implementation and validation of the real-time simulation capabilities needed to represent machine dynamics. Similarly, work on autonomous heavy vehicles, such as Fu et al. [21], highlights the growing role of high-fidelity simulation for training and validating control algorithms. While important for accurate machine-level modeling, these approaches are often task-specific and rely on computationally demanding physics engines, which complicates their integration into a timely DT architecture for operational decision support.

In summary, the literature substantiates the promise of DTs for heavy, off-road industries, but also reveals a recurring systems challenge: existing approaches either provide high-level, largely static frameworks or focus on computationally intensive, low-level machine models. A key gap remains in developing a domain-specific framework that integrates multi-scale models into a computationally efficient architecture capable of delivering real-time, interpretable decision support across quarry operations.

2.3 Modeling and Optimization for Quarry Digital Twins

Quarrying comprises complex workflows involving many interdependent activities [5]. Simulation is therefore widely used to describe site behavior and to test operational strategies. At the same time, a simulation of an entire site is difficult to run at high detail because representing machines, tasks, and their interactions can be computationally expensive. This leads to a trade-off between model fidelity and feasible runtime [36]. In this context, digital twin (DT) developments often combine simulation, optimization, and data-driven components so that site-level decisions can be explored with reasonable latency while still reflecting relevant physical and operational constraints.

Reviews and mapping studies summarize this development. Xu et al. [61] provide a systematic review of construction site-layout planning algorithms, presenting an overview of optimization methods for site arrangements. Guo and Zhang [23] survey multi-objective optimization in construction project management and discuss trends toward multi-criteria decision support and more data-oriented tools. Fernandes et al. [19] categorize approaches to the optimization of earthworks-planning and note the importance of representing real-world interferences and constraints. Together, these works describe a shift from more isolated optimization formulations toward approaches that are increasingly combined with simulation to better reflect site operations.

Alongside these general frameworks, many studies focus on simulation-based optimization for specific planning problems. Burdett and Kozan [10] presented an integrated operations approach to optimize earthwork allocation, sequencing, and routing in large projects, treating material movement as a combined scheduling and routing problem. Kaveh and Vazirinia [33] present a comparative analysis of four metaheuristic algorithms for site layout planning formulated as a quadratic assignment problem, and report that enhanced VPS (EVPS) performs well compared to charged system search, whale optimization, and standard VPS. These approaches support the evaluation of many candidate layouts and haul routes that would be difficult to test manually. Related work has also incorporated spatial technologies (e.g., GPS and GIS) to improve how spatial conditions are represented; Montaser et al. [39] showed that including geospatial data can improve the field estimation of productivity.

A recurring theme in this literature is the need to connect models across levels of detail, from site processes down to equipment behavior. While multi-scale integration in quarry DTs is still commonly reported as difficult in practice and many implementations remain confined to isolated levels [44, 51, 56], earlier work on hybrid and multi-level simulation provides useful reference points for how such coupling could be organized. Discrete-event simulation (DES) and agent-based modeling (ABM) offer different abstractions for site-level behavior: DES focuses on process flows, queues, and resource constraints, whereas ABM emphasizes autonomous machines and operators with local decision rules whose interactions give rise to emergent system dynamics. The choice of DES, ABM, or a hybrid formulation influences how interfaces to continuous machine-dynamics models are defined [34, 62]. Shuang and Xifan [54] introduced a combined discrete-event and continuous simulation model for a hybrid manufacturing system to capture both high-level flow and lower-level dynamics, and Delbrügger et al. [16] proposed a multilevel simulation concept for production systems that links analyses across resolution levels. In a quarry DT setting, these ideas suggest possible ways to structure interfaces between planning, operations, and machine dynamics, but adapting them to site-wide quarry contexts remains non-trivial due to differences in environments, data, and operational variability. Overall, DES tends to be preferable when process flows, queues, and capacity constraints are the primary focus and transparency and computational efficiency are critical, whereas ABM becomes more attractive when heterogeneous actors, local decision rules, and human–machine interactions need to be captured explicitly; in practice, the choice is driven by the dominant decision questions, available data for calibration, and acceptable trade-offs between model complexity, interpretability, and runtime.

In parallel, data-driven methods are increasingly used within DT systems to complement simulation and optimization. One role is to reduce the cost of high-fidelity simulation by learning surrogate models for slow components (e.g., engine dynamics, drivetrain behavior, or vehicle–terrain interaction). Instead of running a physics solver at every step, supervised learning models and hybrid physics–ML approaches can estimate key quantities such as torque, power, and energy consumption from streaming sensor telemetry [11, 12]. When embedded in a DT, such surrogates can support a faster assessment of scenarios and provide timely estimates of machine and operation KPIs for planning and monitoring.

Because these learned components may influence decisions, model accuracy and interpretability are also relevant. Explainable AI (XAI) techniques, including SHAP value attribution and feature selection methods, can help identify which inputs are most informative and reduce unnecessary signals [35, 37]. This can make surrogate models easier to inspect, validate and can support clearer interfaces between DT levels by providing compact, consistent outputs (e.g., predicted fuel use or component stress) that higher-level models can use to update their own states.

Data-driven methods are also used on the control side, where the DT is applied to task-level operational decision-making for activities such as excavation and loading. In off-road operation, the soil–tool interaction and changing material conditions are difficult to model, and performance can depend on perception quality and machine condition. Dadhich et al. [13] identify these as barriers to autonomous earthmoving, and later work showed that neural networks can learn aspects of expert operator behavior, while also indicating that imitation-based policies can be sensitive to shifts in operating conditions [13, 14].

More recent work explores reinforcement learning (RL) as another data-driven approach for task-level control [38]. RL learns a policy from interaction with an environment using reward signals that reflect objectives such as productivity and energy use. Model-free RL has been demonstrated in high-fidelity simulators and controlled trials for off-road equipment, including loading tasks for wheel loaders and load–haul–dump machines [6, 7, 31], and has been reported to outperform engineered baselines in the same simulated settings [17, 49]. A practical issue is the cost of interaction, as model-free RL can require many training episodes [55]. To reduce reliance on repeated expensive simulation, model-based RL approaches learn approximate dynamics “world models” and use them to generate imagined rollouts for policy improvement [28]. Within a DT, this can support faster what-if exploration, followed by validation in higher-fidelity simulation, and more cautious transfer to real machines.

In general, the literature suggests that simulation-based optimization and data-driven methods are often treated as complementary within the development of quarry DT. Simulation and optimization provide structured representations of site processes and constraints, while data-driven components can support faster prediction, clearer model interfaces, and learning-based control concepts that can be tested in virtual environments before being considered for field use.

2.4 Summary and Research Direction

The preceding sections reveal a characteristic trade-off in digital twin research for off-road operations: existing approaches tend to address either strategic site planning or detailed machine dynamics, but rarely both within a unified architecture [9, 21, 43]. High-level frameworks provide useful conceptual guidance, but remain detached from the physics of machine–material interaction [53, 56]. In contrast, high-fidelity vehicle simulations capture transient dynamics with precision but are too computationally expensive for repeated evaluation in operational scenarios [21]. Data-driven models—whether surrogate predictors or learning-based controllers—have demonstrated value in isolation, but are rarely integrated into a multi-level structure where their outputs inform decisions at other scales [11, 45].

A practical digital twin in the quarry site domain should be able to support decisions across multiple time horizons: strategic planning, operational scheduling, and machine-level control. Each horizon imposes different requirements on the fidelity, latency, and type of output needed. Current solutions rarely make these requirements explicit or demonstrate how models at one level supply inputs to another in a timely and consistent manner.

These observations motivate the direction of this thesis. Rather than pursuing a monolithic site model or a single high-fidelity machine simulation, the work focuses on multi-level integration and computational efficiency as enabling principles for digital twins in quarry operations—coupling site, operation, and machine models through defined interfaces, with surrogate predictors and learning-based controllers embedded as functional components rather than isolated tools. The complete integration with live data and field deployment lies beyond the current scope; the emphasis here is on architectural foundations and representative data-driven components, whose formulation and evaluation are detailed in Section 6.1 and the associated publications.

Chapter 3

Research Methodology

3.1 Design science research

Design Science Research (DSR) is a problem-solving research paradigm that creates and evaluates artifact models, methods, and systems to address real needs while contributing to scientific knowledge [30]. It balances relevance (solving the right problem) and rigor (grounding in the knowledge base) through iterative build–evaluate cycles.

In Hevner's three-cycle view, the relevance cycle draws requirements from the application environment, the design cycle iteratively builds and evaluates artifacts, and the rigor cycle grounds design in the scientific knowledge base. Figure 3.1 places this thesis in that framing: the environment includes practitioners, organizations, and technologies; the knowledge base includes foundations and methods; the design cycle produces and evaluates artifacts (requirements mapping, multilevel framework, machine-level data-driven components).

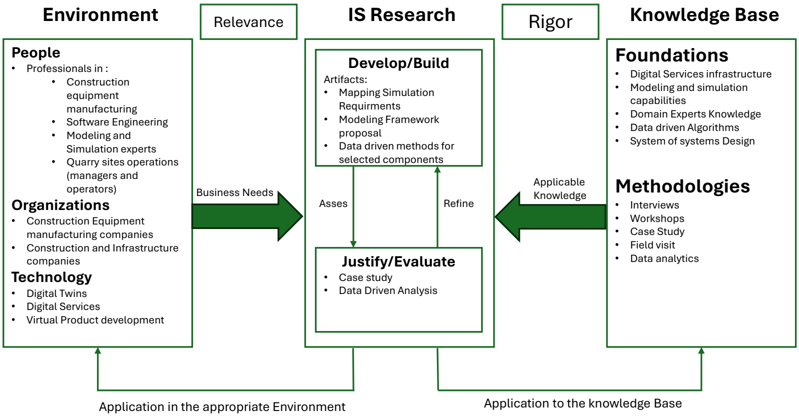

Figure 3.1: The design-science research paradigm employed during the research, adjusted to our context [30].

To operationalize this paradigm, we follow the Design Science Research Methodology (DSRM) of Peffers et al. [42], which structures design science into six iterative activities. Figure 3.2 outlines these activities and indicates where this thesis enters the process.

In terms of Figure 3.1, the environment (people, organizations, technology) drives problem identification and objective setting through interviews, workshops, and site visits (relevance). The knowledge base (foundations, methods) informs the choice of KPIs, modeling approaches, and artifact design

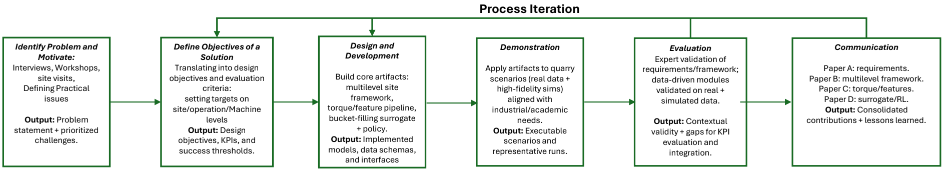

Figure 3.2: Process level (DSRM): six steps and their instantiation in our work; executed iteratively with feedback loops [42].

(rigor). The artifacts reported in this thesis were then built and demonstrated using data and simulator contexts from the environment; they were evaluated—feeding refinements back into design and contributing justified design knowledge to the knowledge base—and the outcomes were communicated (Papers A–D). Taken together, Figure 3.2 provides the process view of Figure 3.1.

Digital twins for quarry sites must be customized to each site's unique terrain, material properties, fleet composition, and regulatory requirements, while continuously adapting to operational changes like equipment upgrades or environmental shifts. DSR provides the paradigm for such context-dependent design; DSRM supplies the process discipline to move from needs to evaluated artifacts. Table 3.1 summarizes DSRM activities and supporting evidence via included papers

Table 3.1: DSRM steps, thesis activities, and evidence.

DSRM step	Activity in this thesis	Evidence
Problem identification	• Interviews, workshops, and site visits identify decision-support gaps, constraints, and KPIs in quarry operations.	Paper A
Define objectives	• Requirements and modeling targets consolidated across site, operational, and dynamics levels.	Paper A
Design & development	• Multilevel Modeling framework • Machine-level components: – Efficient torque surrogate – World-model for RL bucket filling	Papers B, C and D
Demonstration	• Component-level demonstrations: – Torque model – RL policy in high-fidelity simulation	Papers C and D
Evaluation	• Expert reviews • Feasibility of interfaces and data flow • R^2/RMSE/MAPE + SHAP on time-ordered splits • Back-in-sim cycle-time validation	Papers A, B, C and D
Communication	• Peer-reviewed publications and iterative stakeholder reviews with the industrial partner	Papers A, B, C and D

3.2 Research Setting

This licentiate thesis was conducted in collaboration with Volvo Construction Equipment AB (Volvo CE) and within two joint Vinnova-funded projects: TRUST-SOS (TRUSTed Site Optimization Solutions) and TESTED-SOS (Tested Site Optimization Solutions) [3, 4]. TRUST-SOS focused on developing data-driven models and digitalized services to increase the level of trusted decision-making to optimize off-road site systems; the project was coordinated by Volvo CE within the FFI "Fossil-free mobile work machines" venture (Nov. 2021–Dec. 2024) [3].

Building on that foundation, TESTED-SOS advances the work from concept to real-world validation, aiming to improve operational efficiency, reduce downtime, and increase safety at quarry, mining, and construction sites through site-level digital representations, simulation and optimization—thereby enabling new service models and integrated transport solutions [4].

The academic–industrial collaboration is organized through INDTECH – the Industrial Technology Graduate School at Mälardalen University, funded by the Knowledge Foundation (KK-stiftelsen) [1]. INDTECH provides an Industry 4.0 and applied-AI training and research environment with embedded industrial PhD projects, linking IoT/ML/optimization research to concrete industrial needs; this structure supports the thesis's co-creation with Volvo CE.

3.3 Synthesis of Contributions: Digital Twin Framework and Dynamic Components for Quarry Operations

This section explains how the four papers connect. It outlines the progression and rationale behind each study, showing how, together, they contribute to developing a digital twin modeling framework for quarry operations and identify selected components that align with the dynamic level proposed in the framework.

The work progresses from identifying industrial requirements, to proposing a foundational modeling framework, which establishes layer definitions and interfaces for quarry digital twins, with full implementation deferred to future work [22], and finally to developing data-driven components that address computational challenges.

Paper A establishes the baseline by mapping, through interviews and workshops, what a quarry-scale digital twin must support at three levels: site (KPIs such as productivity, cost, energy, and regulatory limits), operations (task planning, coordination, and resource allocation), and dynamics (machine behavior, material interaction, and real-time control). The study also highlights two recurring issues: models must be modular so components can be added, changed and linked, and there is a constant trade-off between accuracy and computational feasibility. These points set the scope and justify a structured, multilevel approach.

Paper B proposes the multilevel framework. It defines interfaces and data flows between three levels to support multi-resolution model integration. At Level 1 (site), stateflow/discrete-event models handle KPIs, schedules, and mission assignment. These outputs create requirements and bounds for Level 2 (operational), where behavioral models specify and monitor tasks (e.g., loading, hauling) and provide reports and risk signals upward. Level 3 (dynamics) contains machine, geospatial, and geometric models that execute tasks and return performance and state data. Downward flows carry targets, plans, and task commands; upward flows carry feedback and outcomes. Expert workshops indicate the structure matches real decision loops. The remaining challenge is clear: Level 3 components must be fast enough to respond to Level 2 specifications and inform Level 1 KPIs.

Paper C complements the architecture with a component-level surrogate for torque prediction. It combines domain-informed feature pre-selection with SHAP-based ranking to obtain a compact input set and transparent attributions. Tabular learners like (e.g., gradient boosting and SVR) were used because the engine signals are low-dimensional yet nonlinear and nonstationary; these models provide high accuracy with millisecond-scale inference and stable behavior under feature reduction. SHAP explanations make the surrogate contributions legible to engineers and allow tracing predictions to controllable variables. As such, the surrogate can be embedded at Level 3 without becoming a computational bottleneck and with inputs/outputs that clearly map to operational logic and site-level KPIs identified earlier in Paper A.

Paper D addresses dynamics level control by learning a bucket-filling policy via a "world model." An LSTM surrogate is used to capture temporal dependencies and partial observability inherent in tool–material interaction; reinforcement learning is then carried out inside this surrogate to improve sample efficiency and avoid unsafe exploration on the high-fidelity simulator. The approach yields faster policy training while preserving task-level fidelity sufficient for evaluating productivity and energy effects against a baseline controller. This demonstrates an example of dynamic level components that can be responsive enough with operational scheduling at higher levels, while acknowledging that the evidence is limited to simulator-based evaluation under controlled conditions. Also it demonstrates reinforcement learning as a important enabler for tackling complex dynamics-level problems in excavation and loading, and establishes a foundation for field validation. Future exploration of RL capabilities on site simulations and real machinery remains open to exploration.

Together, the four papers form a coherent progression: they elicit quarry digital twin requirements across site, operational, and dynamics levels, define how models should interface across these levels, and develop dynamics-level methods that are computationally efficient and reliable for integration into the proposed framework.

Figure 3.3 summarizes the progression: Paper A defines quarry requirements; Paper B specifies the multilevel framework; Papers C (prediction) and D (control) plug into Level 3.

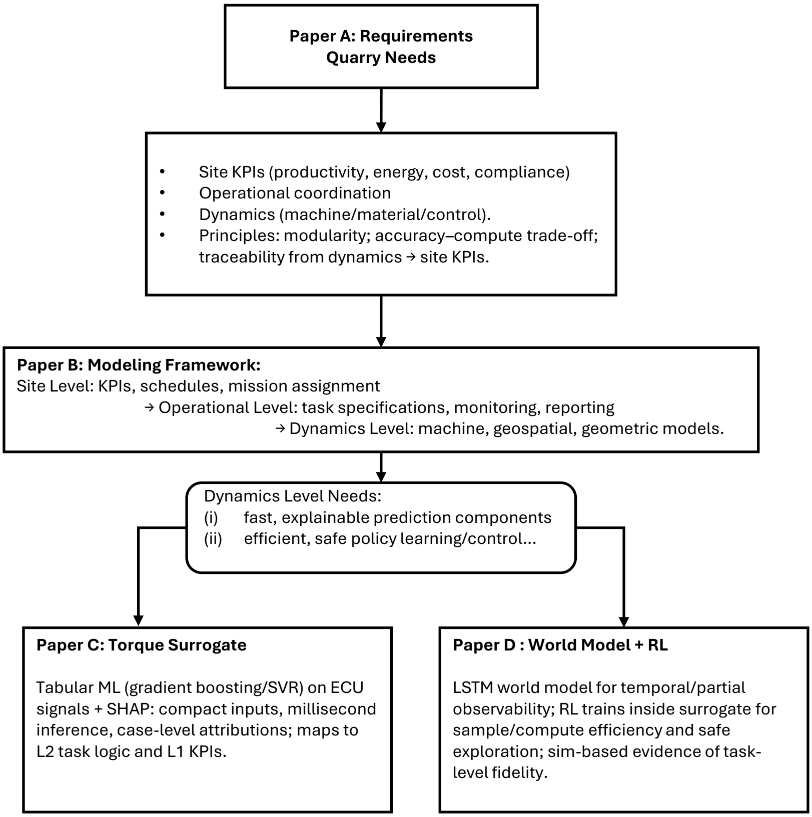

Figure 3.3: Connection of thesis papers.

3.4 Research Process and Design

Within the aforementioned research settings, a broad set of participants (senior researchers, senior software engineers, and site managers) were engaged through semi-structured interviews and facilitated workshops to obtain requirements, operational insight bottlenecks, and co-define use cases for a quarry site digital twin [25, 26]. Guided by these inputs, the research was carried out in three main phases, each corresponding to DSRM cycles and resulting in specific contributions.

1. **Requirements Analysis:** To ground the research in the "Relevance Cycle", this phase focused on capturing industrial reality. Semi-structured interviews and workshops were utilized as primary data collection methods, engaging domain experts to co-define use cases, challenges, and validation criteria. The objective was not to build a model, but to construct a requirements map "the artifact" that established the variables for site, operational, and dynamics levels. This phase answered the "what to build" question and provided validation criteria for subsequent phases. This work is presented in Paper A [25].

2. **Solution Design:** Building on the requirements from Phase 1, this phase moved into the "Design Cycle." The method involved co-design sessions with domain experts to translate the requirements into a coherent modular architecture spanning site-level scheduling, workflow-level behaviour models, and machine-dynamics components, aligned with Volvo CE's digital-transformation objectives. The focus was on defining interfaces—specifically how data flows between high-level discrete event models and low-level continuous dynamics—addressing the integration gap identified in the literature. This work is presented in Paper B [26].

3. **Data-Driven Components:** With the framework establishing the need for fast dynamic models and intelligent control, this phase developed and validated components using data and expert knowledge. The work proceeded in two sub-phases:

    a. **Surrogate prediction Component:** Telemetry data from a Volvo L180 wheel loader were utilized to develop a torque prediction surrogate. The design focus was on feature selection and explainability (SHAP), testing whether expert-guided feature reduction could maintain high fidelity while reducing computational load. This addressed the framework's requirement for computationally efficient machine-dynamics models. This work is presented in Paper C [27].

    b. **RL-based Control Component:** Training data for an autonomous bucket-filling agent were generated using high-fidelity simulation rather than physical telemetry. The method relied on Model-Based Reinforcement Learning using an LSTM world model. This sub-phase demonstrated that data-driven components could not only predict machine behavior but optimize it, addressing the control aspect of the machine-dynamics level. This work is presented in Paper D [40].

Implementation of the complete end-to-end framework constitutes ongoing doctoral research and is therefore outside the scope of this licentiate thesis. The present work delivers the conceptual design, the requirement specification, and initial methodological contributions that enable a future, deployable Digital Twin. Throughout, continuous dialogue with OEM engineers and site practitioners ensured that the proposed abstractions remained grounded in operational realities.

3.5 Data Collection

This research employed multiple data collection strategies, including qualitative stakeholder engagement, quantitative field measurements, and high-fidelity simulations, each aligned to specific research objectives [15]. These complementary data sources underpin the multi-level digital twin framework, integrating qualitative insights, field measurements, and simulation data aligned with their respective validation and development goals.

3.5.1 Requirements and Framework Development

Data for requirements mapping and framework design (Papers A and B) were obtained through iterative semi-structured interviews and expert workshops with Volvo Construction Equipment AB, industrial partners, and academic collaborators. These sessions captured requirements across site, operation, and machine levels by engaging domain experts and stakeholders to co-define use cases, challenges and validation criteria, ensuring the framework reflected practical industrial needs.

3.5.2 Machine-Level Data-Driven Component

Operational measurements for the surrogate torque prediction model (Paper C) were collected from a Volvo L180 wheel loader in a test track facility [59]. The dataset included about 100,000 data points with 37 features, refined through expert curation to 25 relevant variables. This real-world data support model development robust to field variability and operational complexity.

3.5.3 Reinforcement Learning Component

Autonomous control policy data (Paper D) were generated in Volvo's high-fidelity simulation model through automated bucket-filling episodes with systematically varied parameters, resulting in around 120,000 transitions. Controlled perturbations ensured diverse and representative scenarios for reinforcement learning, benefiting from full state observability and repeatability in a realistic yet computationally efficient environment.

3.6 Evaluation Strategy and Validity

Consistent with DSR guidance, evaluation activities were distributed throughout the research process to reduce risk and sharpen design decisions. The evaluation emphasized four aspects: problem framing (assessed through importance and feasibility), design quality (simplicity and consistency), instantiation fidelity (robustness), and solution-in-use effects (effectiveness and efficiency) [42]. Table 3.2 summarizes which metrics and evaluation modes were foregrounded for each artifact.

Table 3.2: Evaluation in this thesis (as reported through the papers).

Artifact (paper)	**Evaluation focus**
Requirements (A)	Expert elicitation and consolidation of site/operations/dynamics requirements; acceptance via workshops/interviews.
Multilevel framework (B)	Feasibility of decision/operational loops and interfaces via expert workshops/interviews.
Torque surrogate (C)	Regression performance (R^2, MSE); interpretability via SHAP.
World model & RL policy (D)	Surrogate fidelity (R^2, NMAE); policy vs. baseline controller on productivity and specific energy.

Chapter 4

Discussion

This chapter examines the implications of the research findings from both theoretical and practical perspectives. The theoretical implications consider how this work contributes to the broader understanding of digital twins in industrial contexts. The practical implications discuss architectural and methodological insights that contribute to the digital twin development of quarry sites.

4.1 Theoretical Implications

This licentiate thesis contributes to the understanding of digital twin frameworks for complex industrial systems, specifically in quarry contexts. The research addresses interconnected challenges that have limited the practical implementation of digital twin in multi-scale industrial environments.

At the architectural level, the proposed three-layer modeling framework provides a structured approach to integrate models across the temporal and spatial scales inherent in quarry operations [26]. By establishing clear interfaces between site, operations and continuous machine dynamics, the framework addresses a gap in connecting high-level planning with low-level dynamics and control. This contribution extends existing digital twin theories from manufacturing contexts to the specific challenges of quarry operations [24, 32], where environments and variable conditions require different modeling considerations. With these levels and interfaces defined, the next challenge is ensuring that lower-level models can be integrated with modeling complexity and computationl effecneicy considerations. With the integration structure in place, the practical bottleneck becomes computational: lower-level models must remain accurate yet efficient enough to run in timely manner.

At the component level, The engine-level surrogate study shows how a critical subsystem can be captured with a compact set of signals while preserving predictive fidelity [27]. Building on requirements elicited earlier in the thesis [25], it demonstrates a concrete procedure for combining expert-guided feature pre-selection with SHAP-based ranking to identify a minimal, interpretable feature set for torque prediction. Methodologically, this contributes a reusable pattern for designing surrogate models in industrial DTs: start from domain knowledge, quantify importance with explainable AI, and converge to the smallest feature subset that still respects the integration and performance needs of the overall architecture. Once lightweight dynamics models have been studied and validated, they can enable fast rollouts for data-driven decision-making and control. This same focus on efficient dynamic representations carries into learned world models for model-based RL, where an explicit speed–fidelity trade-off helps make simulation-based policy learning practical

The LSTM-based world model approach demonstrates one method for applying model-based reinforcement learning in heavy equipment control [40]. By showing that learned dynamics models can reduce training time while maintaining sufficient fidelity for control tasks, this research provides evidence that simulation-based policy learning can be practical for such industrial applications. As quarry sites evolve toward greater autonomy, reinforcement learning is expected to become instrumental in enabling adaptive decision-making across operational scales, from individual machine optimization to coordinated fleet management. This integration has important implications for the capabilities of digital twins, as RL agents can continuously refine their understanding of system dynamics through interaction, potentially transforming digital twins from static simulation tools into adaptive, learning-enabled platforms that improve their predictive accuracy and optimization strategies over time. These contributions offer methodological insights that may inform digital twin development in other industrial domains facing similar challenges of multi-scale integration, computational constraints, and automation requirements.

4.2 Practical Implications

Translating the theoretical contributions into practice, this thesis proposes a multi-level digital twin modeling framework for quarry operations, with demonstrated machine-dynamics components for prediction and control, addressing operational challenges identified through industry collaboration.

The requirements mapping and hierarchical framework provide quarry operators with a structured approach to digital twin adoption. Collaborative stakeholder engagement revealed inter-dependencies between operational levels and clarified how local decisions ripple through site performance. By identifying specific KPIs and interface requirements at each level, the framework establishes clear communication channels and decision-making pathways from strategic planning through operational strategy to machine selection and control. The modular architecture supports incremental implementation, allowing sites to adopt specific components without requiring complete system overhaul. Within this modular setup, a key practical constraint is keeping model components light enough for real-time use.

Balancing model fidelity with real-time performance requires identifying the smallest subset of variables that still capture essential behavior. In this work, expert judgment combined with SHAP-based feature analysis narrowed engine inputs to a handful of signals that accurately reproduce torque responses under diverse operating scenarios. By demonstrating that these core parameters suffice to mirror full-scale dynamics, the surrogate model remains lightweight and responsive, enabling seamless integration into the digital twin without sacrificing computational efficiency. Beyond prediction, these efficient models also make it feasible to evaluate and improve control strategies in simulation at scale.

The reinforcement learning policy achieved 89% productivity gains, 56% energy efficiency improvements, and 50% reduction in cycle time compared to the rule-based driver model in high-fidelity simulation. Before field deployment, validation across varying soil types and terrain conditions is essential; however, these results indicate that reinforcement learning can discover control strategies that are notably different from—and superior to—conventional approaches, which justify further investigation of autonomous control in quarry equipment.

A key consideration is to leverage the digital services already deployed for machine monitoring and health diagnostics rather than introducing standalone platforms. Defining common data interfaces would allow the proposed multilevel framework to ingest real-time telemetry from existing systems, enabling future extensions such as predictive simulations and consequences analyses layered onto current dashboards. By planning integration at the data-exchange level, operators could incrementally enhance their monitoring infrastructure with digital twin capabilities without a major replacement of established tools.

Chapter 5

Conclusions and Future Work

This chapter synthesizes the research contributions by revisiting the initial research questions and outlining directions for future development. The findings establish a foundation for digital twin implementation in quarry operations, while recognizing the ongoing journey from proof-of-concept to full industrial deployment.

5.1 Conclusions

This licentiate thesis addressed two primary research questions through four interconnected studies, contributing on both frameworks and practical components to the field of digital twins for quarry operations.

RQ1: How can a digital twin framework be designed to integrate quarry data and models for real-time, interpretable decision support?

The research answers RQ1 by proposing a three-layer modeling framework supported by stakeholder-driven interface definitions. Paper A mapped requirements through iterative interviews and workshops, revealing important variables and KPIs at the site, operational, and machine levels. Paper B translated these requirements into a modular architecture with defined interfaces that enable bidirectional data exchange between the discrete site model and continuous machine dynamics. By embedding interface specifications into the multilevel framework, the design proposes plug-and-play deployment of individual components while preserving traceable and cohesive decision pathways across the framework hierarchies.

RQ2: How can data-driven approaches be used to model machine-level dynamics within quarry digital twins, while balancing computational efficiency, predictive fidelity, and interpretability for prediction and control?

To address RQ2, this research examined two representative machine-level components. Paper C demonstrated that combining expert-guided feature selection with SHAP analysis yields a minimal feature set (top 5 features) that reliably predicts torque dynamics while maintaining prediction accuracy. This balance of domain insight and statistical validation supports efficient and interpretable digital twin components suitable for quarry operations. Paper D showed that LSTM-based world models facilitate model-based reinforcement learning for bucket-filling control, achieving 89% productivity gains and 56% energy efficiency improvements. These findings confirm that learned dynamics models can support efficient policy development in industrial simulation contexts, with potential for real-world validation.

5.2 Future Work

A primary challenge in digital twin development is balancing model fidelity with computational efficiency for operational timescales relevant to quarry decision-making. Future work shall prioritize physics-informed, data-driven modeling that captures site-specific conditions—soil properties, terrain variability, weather—while maintaining tractable computation. Such hybrid approaches have shown promise in component modeling and warrant investigation for framework refinement.

Another area requiring further investigation is the choice of simulation backbone for site-level. The choice can be discrete-event (focusing on process flows, queues, and capacity constraints), agent-oriented (emphasizing heterogeneous machines and operators with local decision rules), or a combined formulation that blends both perspectives. Selecting between these options requires balancing computational scalability, flexibility to extend the model with new equipment types and layouts, and the ability to represent different operational objectives and control logic, and these tradeoffs should be evaluated under quarry-specific conditions and optimization targets.

Full validation of the proposed framework requires systematic field testing across multiple quarry sites with diverse operational contexts (material types, equipment fleets, and production scales). Validation should assess: (i) model accuracy under real site conditions, (ii) user interface for site planners and operators, (iii) optimization performance on actual production schedules, and (iv) transferability between quarry types and geographies. This evaluation will establish whether the framework can reliably support operational decisions in the quarrying and similar off-road extraction industries.

Chapter 6

Publications

6.1 Summary of included papers

This section summarizes Papers A–C, presenting their abstracts, publication status, and clarifying my contributions and those of my co-authors.

6.1.1 Paper A

Title: Mapping simulation optimization requirements for construction sites: A study in the heavy-duty vehicles industry

Authors: Abdulkarim Habbab, Anas Fattouh, Bobbie Frank, Koteshwar Chirumalla, Markus Bohlin

Venue: The 64th International Conference of Scandinavian Simulation Society (SIMS2023), Västerås, Sweden. doi: 10.3384/ecp200047,

Abstract: The construction and mining industry comprises complex operations and interactions between various actors at different levels. Simulation has emerged as a valuable tool in this domain to better understand the site's behavior and optimize its operation. However, developing a simulation platform that can handle all the operations on the site is challenging due to the computational cost of the digital representation of reality along with the required accuracy level. This paper aims at extracting and mapping the optimization requirements of construction sites at three main levels: site level, operational level and dynamics level. More precisely, this work seeks to define and map the most important requirements between these levels that ensure simulation credibility and reliability. Based on interviews with experts in the domain, both from academia and industry, several key insights and recommendations emerged: at the site level, the layout and the key performance indicators, such as productivity, time, cost, number of machines and workers, need to be modeled and simulated. At the operational level, the simulation platform must include the main activities, such as loading, excavating, transporting and dumping. Moreover, the dynamics level should involve machine models and their interactions with the site's environment, such as earthmoving, drilling, excavating and blasting.

Status: Published

My contribution: I was the main author of the paper and presented it at the SIMS2023 (September) in Västerås, Sweden.

Co-authors' contribution: My co-authors gave important feedback on the manuscript, contributing to its final version.

6.1.2 Paper B

Title: A Multilevel Modelling Framework for Quarry Site Operations

Authors: Abdulkarim Habbab, Anas Fattouh, Bobbie Frank,Elianne Lindmark, Koteshwar Chirumalla, Markus Bohlin

Venue: Proceedings of the 12th ACM/IEEE International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems (SESoS '24). Lisbon, Portugal, 2024, p. 61-64.

doi: 10.1145/3643655.3643881

Abstract: Quarry sites are complex systems that involve several heavy machines, equipment, people, and management systems working together in an unstructured off-road environment. Gaining accurate insights about these sites requires integrating models at different levels to enable a holistic view systems and processes involved and facilitate effective planning, coordination, and decision-making. In this paper, a multi-level modelling framework is proposed to provide an overall structure for the modelling of quarry sites. The motivation for this framework is drawn from insights gained through a large manufacturing company in the heavy-duty vehicle industry, providing a practical perspective on the modeling approach. The framework integrates models of different operations on site enabling effective simulation and optimization and leading to better understanding of the workflow on site and pointing out any possible bottlenecks. The feasibility of the proposed framework was validated through workshops that included a panel of experts in different areas of the field of off-road machinery production company.

Status: Published

My contribution: I was the main author of the paper and presented it at the SESoS '24 (April) in Lisbon, Portugal.

Co-authors' contribution: My co-authors gave important feedback on the manuscript, contributing to its final version.

6.1.3 Paper C

Title: Efficient Torque Prediction for Digital Twins in Quarry Operations: A Data-Driven and Expert-Guided Approach

Authors: Abdulkarim Habbab, Anas Fattouh, Bobbie Frank,Mohammad Loni, Koteshwar Chirumalla, Markus Bohlin

Venue: The IEEE International Conference on Industrial Informatics, KunMing, China 2025, pp. 1-8. doi: 10.1109/INDIN64977.2025.11279455,

Abstract: Quarry sites present unique operational challenges where the performance of heavy machinery is critical for maintaining efficiency and safety. In such environments, accurate torque prediction is essential for effective engine management and optimal task execution. This work addresses the torque prediction challenge for a wheel loader operating in quarry conditions by proposing a structured three-phase approach to feature selection that reduces model complexity while preserving predictive accuracy. In the first phase, features are selected based on domain expertise to capture the physical and operational realities of quarry machinery. A comprehensive set of features is then employed to establish a robust performance baseline. In the final phase, a data-driven analysis using SHapley Additive Explanations (SHAP) identifies the top five features that most significantly impact torque prediction. Model efficacy was validated via cross-validation, with R-squared and mean-squared error serving as the key performance indicators. Comparative analysis reveals that while SHAP-ranked features yield statistically optimal results, the expert-selected features are more aligned with the practical requirements of quarry operations.These findings support the design of efficient, interpretable digital twins for real-time decisions in challenging environments.

Status: Published

My contribution: I was the main author of the paper and presented it at the INDIN (July) in Kunming, China.

Co-authors' contribution: My co-authors gave important feedback on the manuscript, contributing to its final version.

6.1.4 Paper D

Title: Reinforcement Learning with World Models for Autonomous Excavation Optimization in Wheel Loaders

Authors: Duarte Morais, Abdulkarim Habbab, Shehr Bano Fatima, Alexandre Proutiere

Venue: The 66th International Conference of Scandinavian Simulation Society, SIMS 2025, Stavanger, Norway, September 23-24, 2025. doi: 10.1016/j.ifacol.2025.12.184.

Abstract: Automating the bucket-filling task in wheel loaders is challenging due to the complex, nonlinear interaction between the bucket and granular material. This work presents a modelbased reinforcement learning approach to optimize the bucket-filling strategy for Zeux, Volvo's autonomous electric wheel loader concept. A Long Short-Term Memory (LSTM) surrogate model is trained on data from Volvo's high-fidelity simulator to emulate realistic dynamics, enabling efficient policy training using Proximal Policy Optimization (PPO) with imagined rollouts. This reduces computational cost and eliminates the need for direct interaction with the high-fidelity simulator. Compared to Volvo's current rule-based driver model, the learned policy achieves 89% improvement in productivity and 56% increase in energy efficiency. Our results show that world models can accelerate reinforcement learning for heavy machinery control, enabling the discovery of strategies that outperform controllers based on human expert behavior.

Status: Published

My contribution: I was a co-author of the paper where I designed and wrote the methodology and helped setting up the experiment along with reviewing the full paper before submission and updating it after reviewers feedback, and supervising the main author master thesis work

Co-authors' contribution: the other co-authors gave important feedback on the experimental setup, and manuscript, contributing to the final version.

Bibliography

[1] INDTECH – industrial technology graduate school. `https://indtech-graduateschool.se/`, 2021. Industrial Technology Graduate School at Mälardalen University, funded by the Knowledge Foundation (KK-stiftelsen).

[2] Iso 23247-1:2021 — automation systems and integration — digital twin framework for manufacturing — part 1: Overview and general principles, 2021. URL `https://www.iso.org/standard/75066.html`. Accessed: 2026-02-04.

[3] TRUST-SOS – trusted site optimization solutions. `https://www.vinnova.se/p/trust-sos/`, 2021. Vinnova FFI project Ref. 2021-02551, Nov. 2021–Dec. 2024, coordinated by Volvo Construction Equipment AB.

[4] TESTED-SOS – tested site optimization solutions. `https://www.vinnova.se/en/p/tested-site-optimization-solutions-tested-sos/`, 2024. Vinnova FFI project Ref. 2024-03678, Nov. 2024–Sep. 2027, coordinated by Volvo Construction Equipment AB.

[5] Satyajith Amaran, Nikolaos V. Sahinidis, Bikram Sharda, and Scott J. Bury. Simulation optimization: a review of algorithms and applications. *Annals of Operations Research*, May 2016. doi: 10.1007/s10479-015-2019-x.

[6] Osher Azulay and Amir Shapiro. Wheel loader scooping controller using deep reinforcement learning. *IEEE Access*, 9:24145–24154, 2021. doi: 10.1109/ACCESS.2021.3056625.

[7] Sofi Backman, Daniel Lindmark, Kenneth Bodin, Martin Servin, Joakim Mörk, and Håkan Löfgren. Continuous control of an underground loader using deep reinforcement learning. *Machines*, 9(10):216, 2021. doi: 10.3390/machines9100216.

[8] Stephan Baumgart, Joakim Fröberg, and Sasikumar Punnekkat. A state-based extension to stpa for safety-critical system-of-systems. In *2019 4th International Conference on System Reliability and Safety (ICSRS)*, pages 246–254, 2019. doi: 10.1109/ICSRS48664.2019.8987632.

[9] Alessandro Bertoni, Raj Jiten MacHchhar, Tobias Larsson, and Bobbie Frank. Digital Twins of Operational Scenarios in Mining for Design of Customized Product-Service Systems Solutions. In *Procedia CIRP*. Elsevier B.V., 2022. doi: 10.1016/j.procir.2022.05.290.

[10] R.L. Burdett and E. Kozan. An integrated approach for earthwork allocation, sequencing and routing. *European Journal of Operational Research*, 2014. doi: https://doi.org/10.1016/j.ejor.2014.04.036.

[11] Rafael Canal, João Paulo Araujo Bonomo, Rodrigo Santos de Carvalho, and Giovani Gracioli. Torque regression using machine learning techniques in automotive ECUs. *Design Automation for Embedded Systems*, December 2024. doi: 10.1007/s10617-024-09289-9.

[12] Andrea Coraddu, Miltiadis Kalikatzarakis, Gerasimos Theotokatos, Rinze Geertsma, and Luca Oneto. Physical and Data-Driven Models Hybridisation for Modelling the Dynamic State of a Four-Stroke Marine Diesel Engine. In *Energy, Environment, and Sustainability.* Springer Nature, 2022. doi: 10.1007/978-981-16-8618-4_6.

[13] S. Dadhich, U. Bodin, and U. Andersson. Key challenges in automation of earth-moving machines. *Automation in Construction*, August 2016. doi: 10.1016/j.autcon.2016.05.009.

[14] Siddharth Dadhich, Fredrik Sandin, and Ulf Bodin. Predicting bucket-filling control actions of a wheel-loader operator using a neural network ensemble. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6, 2018. doi: 10.1109/IJCNN.2018.8489388.

[15] Shanna Daly, A. McGowan, and Panos Papalambros. Using qualitative research methods in engineering design research. *Proceedings of the International Conference on Engineering Design, ICED*, January 2013.

[16] Tim Delbrügger, Matthias Meißner, Andreas Wirtz, Dirk Biermann, Johanna Myrzik, Jürgen Rossmann, and Petra Wiederkehr. Multi-level simulation concept for multidisciplinary analysis and optimization of production systems. *International Journal of Advanced Manufacturing Technology*, August 2019. doi: 10.1007/s00170-019-03722-1.

[17] Pascal Egli, Dominique Gaschen, Simon Kerscher, Dominic Jud, and Marco Hutter. Soil-adaptive excavation using reinforcement learning. *IEEE Robotics and Automation Letters*, 7(4):9778–9785, 2022. doi: 10.1109/LRA.2022.3189834.

[18] Jiyu Feng, Wei Chen, Changyi Liu, Xinhui Liu, Bingwei Cao, and Feixiang Xu. Advancements and challenges in earthmoving equipment automation using core technologies. *Automation in Construction*, 179, November 2025. doi: https://doi.org/10.1016/j.autcon.2025.106472.

[19] Pedro G.P.S. Fernandes, Ernesto F. Nobre Júnior, and Bruno de A. Prata. Optimization of earthworks planning: a systematic mapping study. *Canadian Journal of Civil Engineering*, 2022. doi: https://doi.org/10.1139/cjce-2022-0185.

[20] Jiali Fu, Erik Jenelius, and Haris Koutsopoulos. Optimizing Fleet Selection For Earthmoving Operations. In *Proceedings of the 13th East Asia-Pacific Conference on Structural Engineering and Construction (EASEC-13)*, 2013. doi: 10.3850/978-981-07-5354-2_cpm-15-237.

[21] Jiali Fu, Jan Åslund, and Erik Uhlin. A simulation framework for off-road transport operations using electric construction vehicles. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 2539–2544, 2019. doi: 10.1109/ITSC.2019.8916996.

[22] Shirley Gregor and Alan R. Hevner. Positioning and presenting design science research for maximum impact. *MIS Quarterly*, June 2013. doi: 10.25300/MISQ/2013/37.2.01.

[23] Kai Guo and Limao Zhang. Multi-objective optimization for improved project management: Current status and future directions. *Automation in Construction*, July 2022. doi: https://doi.org/10.1016/j.autcon.2022.104256.

[24] Sebastian Haag and Reiner Anderl. Digital twin – Proof of concept. *Manufacturing Letters*, 15: 64–66, January 2018. ISSN 22138463. doi: 10.1016/j.mfglet.2018.02.006.

[25] Abdulkarim Habbab, Anas Fattouh, Bobbie Frank, Koteshwar Chirumalla, and Markus Bohlin. Mapping simulation optimization requirements for construction sites: A study in the heavy-duty vehicles industry. In *64th International Conference of Scandinavian Simulation Society (SIMS 2023)*. Linköping Electronic Conference Proceedings 192, October 2023. doi: 10.3384/ecp200047.

[26] Abdulkarim Habbab, Anas Fattouh, Bobbie Frank, Elianne Lindmark, Koteshwar Chirumalla, and Markus Bohlin. A Multilevel Modelling Framework for Quarry Site Operations. In *Proceedings - 2024 IEEE/ACM 12th International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems, SESoS 2024*, pages 61–64. Association for Computing Machinery, Inc, 4 2024. ISBN 9798400705571. doi: 10.1145/3643655.3643881.

[27] Abdulkarim Habbab, Anas Fattouh, Mohammad Loni, Koteshwar Chirumalla, Bobbie Frank, and Markus Bohlin. Efficient torque prediction for digital twins in quarry operations: A data-driven and expert-guided approach. In *2025 IEEE 23rd International Conference on Industrial Informatics (INDIN)*, 2025. doi: 10.1109/INDIN64977.2025.11279455.

[28] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *Nature*, April 2025. doi: 10.1038/s41586-025-08744-2.

[29] Mario Hermann, Tobias Pentek, and Boris Otto. Design principles for industrie 4.0 scenarios. In *Proceedings of the Annual Hawaii International Conference on System Sciences*, volume 2016-March, pages 3928–3937. IEEE Computer Society, 3 2016. ISBN 9780769556703. doi: 10.1109/HICSS.2016.488.

[30] Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design science in information systems research. *MIS Quarterly*, March 2004.

[31] Jianfei Huang, Dewen Kong, Guangzong Gao, Xinchun Cheng, and Jinshi Chen. Data-driven reinforcement-learning-based automatic bucket-filling for wheel loaders. *Applied Sciences*, 11(19), October 2021. doi: 10.3390/app11199191.

[32] Haifan Jiang, Shengfeng Qin, Jianlin Fu, Jian Zhang, and Guofu Ding. How to model and implement connections between physical and virtual models for digital twin application. *Journal of Manufacturing Systems*, pages 36–51, January 2021. doi: 10.1016/j.jmsy.2020.05.012.

[33] Ali Kaveh and Yasin Vazirinia. Construction site layout planning problem using metaheuristic algorithms: A comparative study. *Iranian Journal of Science and Technology, Transactions of Civil Engineering*, 2019. doi: 10.1007/s40996-018-0148-6.

[34] Ali Khodabandelu and JeeWoong Park. Agent-based modeling and simulation in construction. *Automation in Construction*, 131, 2021. doi: https://doi.org/10.1016/j.autcon.2021.103882.

[35] Scott M. Lundberg, Gabriel Erion, Hugh Chen, Adam Wright, et al. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, pages 56–67, 2020. doi: 10.1038/s42256-019-0138-9.

[36] Raj Jiten Machchhar. *Changeability Assessment in Complex Systems to Support Early-Stage Design Decisions*. Doctoral dissertation, Blekinge Institute of Technology, Karlskrona, Sweden, September 2025.

[37] Wilson E. Marcilio-Jr and Danilo M. Eler. From explanations to feature selection: assessing shap values as feature selection mechanism. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 340–347, November 2020. doi: 10.1109/SIBGRAPI51738.2020.00053.

[38] Thomas M. Moerland, Joost Broekens, Aske Plaat, and Catholijn M. Jonker. Model-based reinforcement learning: A survey. *Foundations and Trends in Machine Learning*, page 1–118, January 2023. doi: 10.1561/2200000086.

[39] Ali Montaser, Ibrahim Bakry, Adel Alshibani, and Osama Moselhi. Estimating productivity of earthmoving operations using spatial technologies1this paper is one of a selection of papers in this special issue on construction engineering and management. *Canadian Journal of Civil Engineering*, 2012. doi: 10.1139/l2012-059.

[40] Duarte Morais, Abdulkarim Habbab, Shehr Bano Fatima, and Alexandre Proutiere. Reinforcement learning with world models for autonomous excavation optimization in wheel loaders. *IFAC-PapersOnLine*, 2025. doi: https://doi.org/10.1016/j.ifacol.2025.12.184. 66th International Conference of Scandinavian Simulation Society SIMS 2025.

[41] Madara Premawardhana Dassanayake Mudiyanselage and Harin Sellahewa. Digital twins as a framework for iot applications: A review. In *7th SLAAI International Conference on Artificial Intelligence (SLAAI-ICAI)*, 2023. doi: 10.1109/SLAAI-ICAI59257.2023.10365018.

[42] Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of Management Information Systems*, December 2007. doi: 10.2753/MIS0742-1222240302.

[43] M. Pregnolato, S. Gunner, E. Voyagaki, R. De Risi, N. Carhart, G. Gavriel, P. Tully, T. Tryfonas, J. Macdonald, and C. Taylor. Towards Civil Engineering 4.0: Concept, workflow and application of Digital Twins for existing infrastructure. *Automation in Construction*, 2022. doi: 10.1016/j.autcon.2022.104421.

[44] Juncong Qu, M. S. Kizil, Mohsen Yahyaei, and P. F. Knights. Digital twins in the minerals industry – a comprehensive review. *Mining Technology*, 132(4):267–289, 2023. doi: 10.1080/25726668.2023.2257479.

[45] Adil Rasheed, Omer San, and Trond Kvamsdal. Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 2020. doi: 10.1109/ACCESS.2020.2970143.

[46] Kay Rogage, Elham Mahamedi, Ioannis Brilakis, and Mohamad Kassem. Beyond digital shadows: A Digital Twin for monitoring earthwork operation in large infrastructure projects. *AI in Civil Engineering*, December 2022. doi: 10.1007/s43503-022-00009-5.

[47] Farnaz Sadeghpour and Mohsen Andayesh. The constructs of site layout modeling: An overview. *Canadian Journal of Civil Engineering*, January 2015. doi: 10.1139/cjce-2014-0303.

[48] Radhya Sahal, John G. Breslin, and Muhammad Intizar Ali. Big data and stream processing platforms for industry 4.0 requirements mapping for a predictive maintenance use case. *Journal of Manufacturing Systems*, January 2020. doi: https://doi.org/10.1016/j.jmsy.2019.11.004.

[49] Rodrigo Salas, Francisco Leiva, and Javier Ruiz del Solar. Autonomous loading of ore piles with load-haul-dump machines using deep reinforcement learning. *Expert Systems with Applications*, 264:125770, 2025. doi: https://doi.org/10.1016/j.eswa.2024.125770.

[50] Samad M. E. Sepasgozar. Differentiating digital twin from digital shadow: Elucidating a paradigm shift to expedite a smart, sustainable built environment. *Buildings*, 2021. doi: 10.3390/buildings11040151.

[51] Martin Servin, Tomas Berglund, and Samuel Nystedt. A multiscale model of terrain dynamics for real-time earthmoving simulation. *Advanced Modeling and Simulation in Engineering Sciences*, 8:11, May 2021. doi: 10.1186/s40323-021-00196-3.

[52] Martin Servin, Folke Vesterlund, and Erik Wallin. Digital twins with distributed particle simulation for mine-to-mill material tracking. *Minerals*, 11, 2021. doi: 10.3390/min11050524.

[53] Guodong Shao and Moneer Helu. Framework for a digital twin in manufacturing: Scope and requirements. *Manufacturing Letters*, 2020. doi: https://doi.org/10.1016/j.mfglet.2020.04.004.

[54] Liang Shuang and Yao Xifan. Multi-level modeling for hybrid manufacturing systems using arena and MATLAB. In *Proceedings - 2008 International Workshop on Modelling, Simulation and Optimization, WMSO 2008*. IEEE Computer Society, 2008. doi: 10.1109/WMSO.2008.79.

[55] Laura Smith, J Chase Kew, Xue Bin Peng, Sehoon Ha, Jie Tan, and Sergey Levine. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Robotics: Science and Systems (RSS)*, 2023.

[56] Fei Tao, Bin Xiao, Qinglin Qi, Jiangfeng Cheng, and Ping Ji. Digital twin modeling. *Journal of Manufacturing Systems*, 64:372–389, 2022. doi: https://doi.org/10.1016/j.jmsy.2022.06.015.

[57] H. Tawfik and T. Fernando. A simulation environment for construction site planning. In *Proceedings Fifth International Conference on Information Visualisation*, 2001. doi: 10.1109/IV.2001.942059.

[58] Volvo Construction Equipment. Volvo services & solutions — volvo ce europe. `https://www.volvoce.com/europe/en/volvo-services/`. Accessed: 2026-01-23.

[59] Volvo Construction Equipment. *L180H Wheel Loader Product Specifications.* Volvo CE, 2024. URL `https://www.volvoce.com/europe/en/products/wheel-loaders/l180/`. Accessed: 2026-01-23.

[60] Volvo Construction Equipment. Connected solutions — how we drive change, 2025. URL `https://www.volvoce.com/global/en/about-us/how-we-drive-change/connected-solutions/`. Accessed: 2026-01-23.

[61] Maozeng Xu, Zhongya Mei, Siyu Luo, and Yi Tan. Optimization algorithms for construction site layout planning: a systematic literature review. October 2020. doi: 10.1108/ECAM-08-2019-0457.

[62] Emile Zankoul, Hiam Khoury, and Rita Awwad. Evalutation of agent-based and discrete-event simulation for modeling construction earthmoving operations. In Mikko Malaska and Rauno Heikkilä, editors, *Proceedings of the 32nd International Symposium on Automation and Robotics in Construction and Mining (ISARC 2015)*, pages 1–8, Oulu, Finland, June 2015. International Association for Automation and Robotics in Construction (IAARC). ISBN 978-951-758-597-2. doi: 10.22260/ISARC2015/0014.