DOI: 10.21062/mft.2026.010 http://www.journalmt.com

Machine Learning Regression Approaches for Manufacturing Cost and Time Prediction: A Comprehensive Review

Michal Matějka (0009-0007-0763-5438)[1], Milan Dian (0000-0003-2740-7292)[2], Jan Lhota (0000-0002-9726-4461)[1], Theodor Beran (0000-0003-4210-9525)[1], Vojtěch Hlinák (0009-0005-7358-7505)
[1]Faculty of Mechanical Engineering, CTU in Prague, Karlovo náměstí 13, 121 35, Prague 2 – Nové město. Czech Republic. E-mail: michal.matejka@cvut.cz
[2]Faculty of Mechanical Engineering, J. E. Purkyne University in Usti nad Labem. Pasteurova 3334/7, 400 01 Usti nad Labem. Czech Republic. E-mail: milan.dian@ujep.cz

Today, machine learning regression methods are quietly but fundamentally transforming cost and time estimation in manufacturing: from early pricing to labor planning to operational order management. This survey offers a comprehensive map of approaches - from linear models, to tree ensembles (RF, GBM, XGBoost) and shallow neural networks, to multi-target and tensor regressions that can exploit data structure across BOM items and sequences of operations. With an emphasis on SME conditions, we show how to reconcile three often conflicting requirements of practice: accuracy, explainability, and integration into existing data flows (MES/ERP). The paper presents a comparative taxonomy of methods, recommended validation practices (MAE, RMSE, MAPE, R^2 including confidence intervals) and a pragmatic adoption trajectory: from regularized multiple regressions to tree models to multi-output formulations sharing re-presentations across operations. Consolidated findings show that modern learners consistently outperform traditional baselines when supported by careful flag engineering, drift management, and data standardization. As a major research-application contribution, we propose a unified multi-objective framework for simultaneous cost and time prediction that combines domain (queueing/simulation) features with data-driven regression to enable transparent decision making in pricing and capacity planning. The study thus creates a bridge between theory and manufacturing practice and invites the reader to systematically but achievably deploy ML in everyday decision making.

Keywords: Machine Learning, Regression, Cost Estimation, Time Prediction, Manufacturing

1 Introduction

In today's fast-moving manufacturing environment, cost competitiveness is pivotal particularly for small and medium-sized enterprises (SMEs) that must satisfy increasingly customized and volatile demand under tight resource constraints. Globalized supply chains, fluctuating energy and material prices, and shortening product lifecycles intensify pressure on quotation accuracy and schedule reliability. At the same time, customers expect rapid responses and transparent pricing, often for variants the company has never produced before. Against this backdrop, one of the most demanding and time-intensive tasks is accurate cost estimation, especially early in product development when only fragmentary information is available. The challenge is amplified in engineer-to-order and make-to-order settings, where limited reuse of prior routings, uncertain yields, and evolving specifications create a moving target for both cost and lead-time prediction.

In practice, a customer inquiry initiates a pipeline of activities: feasibility assessment, preliminary design, cost estimation, quotation, order confirmation, and production planning. Each stage introduces and transforms information geometry and tolerance hints, tentative routings, supplier quotes, and risk assumptions which must be consolidated into a coherent estimate. In many SMEs this step still relies on spreadsheets and expert judgment. While workable for repeat products, such manual approaches are subjective, hard to standardize or transfer within the firm, time-consuming, and difficult to scale to changing market conditions [1,2,3,4,5]. Knowledge is frequently siloed in a few senior estimators; when they are unavailable, response times degrade and estimates become inconsistent. Moreover, spreadsheet templates often evolve informally, leading to versioning issues, formula drift, and weak auditability problems that surface during post-mortems when planned and actual costs diverge.

Reliable estimates are essential for profitability and competitiveness: underestimation erodes margins, whereas overestimation risks losing orders. To balance speed with accuracy, firms need methods that can ingest partial inputs, quantify uncertainty, and update as new data arrive. Predictive modeling with machine learning (ML) offers a data-driven alternative to static, assumption-based calculations. By learning directly from historical cases, ML can uncover complex, nonlinear links among product attributes, process parameters, and costs even when engineers cannot specify these relations a priori making ML particularly valuable in the early stages when specifications are incomplete [1,6,7]. In practical terms, this means mapping simple descriptors (e.g., bounding-box dimensions, estimated mass, material class, surface finish level, nominal tolerances, process family) to continuous outcomes such as labor hours, energy use, scrap risk, or total unit cost. With incremental retraining, such models can progressively refine estimates as CAD matures and supplier information solidifies.

Evidence across sectors shows that both supervised and hybrid ML approaches can deliver higher accuracy than traditional regression. Algorithms such as SVM, ANN, Gaussian Process Regression (GPR), decision trees, k-nearest neighbors, and ensemble models (e.g., AdaBoost/DT) consistently outperform classical methods [1,2,8,9]. Typical industrial datasets are heterogeneous and modest in size; here, SVMs and tree ensembles tend to handle nonlinearity and interactions with less manual feature crafting than linear models. For example, in the power transformer domain, GPR predicts man-hours a key cost driver more accurately than expert estimates or alternative ML techniques [2]. In shipbuilding and aerospace assembly, ML-based man-hour prediction surpasses multiple regression and expert judgment [3,4,10]. These improvements often stem from the ability to pool weak signals from many small features (geometry proxies, fixture complexity, setup counts) that collectively explain substantial variance. Similar patterns appear in software effort estimation, where ANN and SVM often dominate competing approaches [11]. Although software and manufacturing differ in process physics, the shared lesson is that flexible, data-driven regressors adapt better than rigid parametric forms when task variety is high.

Beyond labor hours, ML has been applied to forecast material consumption, energy usage, procurement costs, and entire cost structures. Within supply chains, intelligent systems such as ANNs, LS-SVM, and MARS support procurement cost estimation and analysis of cost drivers [6]. These systems help buyers reason about supplier price volatility, minimum order quantities, and currency effects, and they enable "what-if" analysis (e.g., material substitutions or alternate sourcing). In additive manufacturing, models exploit 3D geometry and G-code/process instructions to estimate costs from shape, material, and process parameters [12,13]. Here, features may include toolpath length, deposition rate, travel vs. print time, support volume, and thermal cycles signals that are difficult to model analytically but are naturally captured by data-driven regressors. Analogous strategies extend to airframe manufacturing [8], hybrid manufacturing cost frameworks [5], and large infrastructure projects (e.g., metro and light rail) [9], where heterogeneous inputs and multi-stage processes make handcrafted cost equations brittle. Across these domains, ML not only improves point accuracy but also provides uncertainty bands that aid risk-adjusted pricing and contract negotiation.

More advanced techniques are emerging in specialized settings. In wire-arc additive manufacturing, meta-learning schemes that stack base learners (RF, SVR, XGB, KNN) and ridge regression have been proposed for cost estimation and optimization [14,15]. Stacked generalization leverages complementary inductive biases trees for interactions, kernels for smooth nonlinearities, and instance-based methods for local patterns yielding resilient performance across regimes. Deep learning, including 3D CNNs, supports automatic feature recognition and cost prediction in machining [7]. A typical pipeline detects manufacturability features (pockets, bosses, holes), infers setups and tool changes, and regresses cycle time or cost components end-to-end from point clouds or voxelized CAD. In construction, AI methods are being integrated with parametric modeling to improve early-stage estimates [16], where rapid massing studies and option exploration benefit from learned mappings between parametric attributes and cost indices. Collectively, these developments point toward convergent practices across industries: learned surrogates that couple automated feature extraction with regression to deliver real-time cost and time guidance inside design workflows.

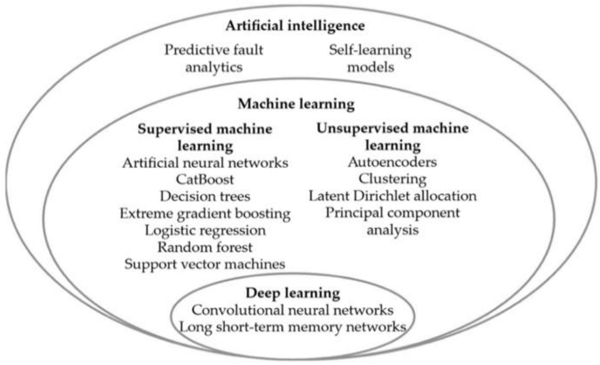

Fig. 1 *Concept of ML Hierachy (Tsallis et al., 2025)*

Despite this promise, important challenges persist. Model performance hinges on data quality and quantity; incomplete or inconsistent datasets undermine reliability [1,13]. SMEs often face sparse historical records, heterogeneous naming conventions, missing timestamps, or partial cost breakdowns issues that necessitate a data-engineering push before ML is viable. Establishing a canonical data model (parts, routings, machines, operations, and standard cost categories), enforcing unit consistency, and instrumenting data capture at quotation, production, and close-out are vital precursors. Cold-start scenarios, new materials, novel routings, or greenfield products require strategies such as transfer learning from adjacent product families, synthetic data augmentation, or conservative priors until sufficient evidence accrues. Interpretability is another concern: while black-box models such as deep ANNs may be accurate, more transparent approaches like MARS or LS-SVM can be preferable for managerial decision-making [5,6]. In practice, explainability tooling (e.g., partial-dependence profiles, SHAP) helps estimators and sales engineers validate drivers, negotiate with customers, and justify price deltas. Finally, robust deployment calls for attention to drift monitoring (supplier changes, machine upgrades), periodic retraining cadences, and governance over who can approve model updates.

This study aims to design and implement a predictive tool for estimating manufacturing cost drivers and pricing orders using a tensor-regression approach enhanced with ML, deployed within the accessible environment of Microsoft Excel. The rationale for Excel is twofold: (i) it is already the lingua franca of many estimation teams, lowering adoption barriers, and (ii) it provides a familiar front-end for data review, feature overrides, and scenario analysis. By integrating historical production and cost data, the tool will predict labor hours, material usage, and energy consumption, narrowing the gap between theoretical estimates and realized performance. The core workflow comprises (1) standardized data ingestion from ERP/MES exports and archived quotations; (2) feature synthesis from early descriptors (dimensions, weight, tolerance class, material family, process route candidates) and, where available, text extraction from RFQs or G-code; (3) model training with cross-validated hyperparameter search and stability checks; and (4) Excel-native scoring and reporting, with guardrails such as confidence thresholds and variance alerts. To aid decision quality, outputs will include point estimates, uncertainty intervals, and sensitivity diagnostics indicating which drivers most influence the estimate. Over time, a lightweight MLOps loop versioned datasets, retraining triggers, and back-testing against actuals will keep the model aligned with changing shop conditions. The outcome is a transition from subjective, manual practices to scalable, data-driven decision-making, strengthening operational efficiency and strategic competitiveness in the Industry 4.0 era.

In sum, modern regression approaches provide SMEs with a pragmatic path to better, faster estimates without demanding “big data” or wholesale system replacement. By starting with disciplined data collection, modest but well-chosen features, and proven learners (SVR, tree ensembles, hybrid stacks), enterprises can raise quotation accuracy, shorten response times, and institutionalize expertise that previously lived only in spreadsheets and memory. As the dataset grows, incremental adoption of deeper models, multi-target formulations, and structured inputs becomes feasible yet the central aim remains unchanged: deliver trustworthy cost and time predictions early and often, where they have the greatest financial leverage.

2 Regression Techniques for Cost and Time Estimation in Manufacturing

Accurate costing and time prediction in manufacturing requires models that handle mixed tabular data (ERP/MES/BOM), tolerate collinearity, capture nonlinear interactions, and remain explainable enough for quotation, planning, and audit. This section reviews the regression families most reported in industrial and adjacent domains, highlighting where each tends to work best and how they have been applied to cost/time targets in practice.

2.1 Multiple Linear Regression (MLR): the starting point

In many firms, multiple linear regression is the first data-driven baseline for mapping engineered features (dimensions, mass, material class, routing indicators) to continuous outcomes such as unit cost, labor hours, or cycle/lead time. Its appeal is transparency and ease of governance; coefficients act as interpretable elasticities for engineering and controller teams. In construction and project costing, classical studies demonstrate MLR’s practicality, but also the limits of strict linearity once interactions and heteroskedastic noise arise common patterns in manufacturing shop-floor data. [17]

In its simplest form, MLR assumes a linear relationship between features x_j and the target y:

$$\hat{y} = \beta_0 + \sum_{j=1}^{p} \beta_j x_j + \varepsilon \tag{1}$$

Its attraction is transparency (coefficients as cost/time elasticities) and straightforward diagnostics. In practice, however, manufacturing data often violate OLS assumptions (non-Gaussian noise, heteroscedasticity, skew). Industrial case work therefore adapts the loss to reflect relative or log-relative error more appropriate for quoting labor hours or setup times and may weight observations to counteract measurement bias across operations with very different lot sizes. These modifications have been shown to improve fit over naïve OLS in factory time-prediction studies, while keeping the model auditable for planners.

As a general rule, MLR is a strong interpretable baseline and remains widely reported across industrial ML surveys; nonetheless, it is frequently outperformed on nonlinear problems by modern ensemble or deep models [17].

However, these single-output regression models assume linear dependence and often ignore more complex non-linear or interrelated influences of inputs. In practice, they also require sufficiently detailed data – product or process parameters – which are only available in later stages of development. In the early stages of product design, especially for customised modifications, only minimal information is available, which causes high uncertainty in estimates. This poses a significant problem for small and medium-sized enterprises (SMEs): without extensive databases of past projects, reliable figures cannot be obtained using traditional analogies and expert estimates. As Hennebold point out, in the engineering SME environment, qualitative methods in the early stages of development are insufficient for reliable economic evaluation and decision-making [18].

2.2 Regularized Linear Regressions

Manufacturing datasets often exhibit multicollinearity (e.g., correlated geometry proxies, overlapping routing flags). Ridge (L2) stabilizes estimates by shrinking coefficients; Lasso (L1) performs embedded feature selection and is useful when many weak descriptors need pruning; Elastic Net blends both to keep correlated groups. Reviews and case syntheses list Lasso among mainstream regressors in industrial analytics; however, on interaction-rich cost panels and early project estimation, tree-boosting families often surpass pure linear models in accuracy reinforcing the role of regularized linear regression as a robust baseline or as a component in stacked ensembles rather than the final word.

Ridge Regression (L2 penalization):

$$min = \sum (y_i - \hat{y}_i)^2 + \lambda \sum \beta_j^2 \tag{2}$$

Lasso Regression (L1 penalization):

$$min = \sum (y_i - \hat{y}_i)^2 + \lambda \sum |\beta_j| \tag{3}$$

More advanced multi-output models apply various regularizations to the W matrix (or, more generally, to the parameter tensor) in order to explicitly exploit inter-target correlations. Low-rank assumptions are often used i.e., that the W matrix can be approximated by the product of smaller matrices, which means that individual targets share common latent factors. Other approaches introduce sparsity, where only a subset of predictors influences groups of outputs, or structured regularization encoding hierarchy or group ties between outputs. The advantage of algorithmic adaptation is that target correlations are modeled directly within a single optimization problem, which often leads to better interpretability (e.g., we can see the structure of the weight W) and sometimes even better predictive accuracy than with distributed models [17,19,20].

2.3 Non-Linear Regression: Decision Trees & Random Forests

Decision Trees (DT) partition the feature space into regions and predict a constant (leaf mean) per region; they capture nonlinearities and interactions without heavy feature engineering but overfit when grown deep [22]. Random Forests (RF) reduce this variance by averaging many de-correlated trees trained on bootstrap samples with feature subsampling. The RF regressor predicts

$$\hat{y}(x) = \frac{1}{T} + \sum_{t=1}^{T} f_t(x), \tag{4}$$

Where:

f_t…The t-th tree and

T…The number of trees.

RFs work well on heterogeneous tabular data (ERP/MES/BOM), require minimal scaling, and provide variable-importance diagnostics useful for identifying dominant cost/time drivers [17,23,24].

Gradient Boosted Machines (GBM) build an additive model by fitting shallow trees to residuals:

$$F_m(x) = F_{m-1}(x) + \nu . h_m(x), \tag{5}$$

With learning rate ν and weak learner h_m chosen to minimize the empirical loss at stage m (e.g., squared error) [24]. A regularized objective highlights the trade-off:

$$h_m = \underset{h}{\operatorname{argmin}} \sum_{i=1}^{n} (y_i - F_{m-1}(x_i) - \nu h_{x_i})^2 + \lambda \|h\|_2^2 , \tag{6}$$

XGBoost implements second-order optimization and explicit regularization/handling of missing values, delivering state-of-the-art accuracy on structured cost/time data with strong robustness to sparsity [25]. In manufacturing and project costing, boosting (including XGBoost/LightGBM) often outperforms linear baselines and RF when interactions and mild nonlinearities dominate [17,27].

Multi-output capability. Trees and forests can emit vector leaves to perform multi-output regression (e.g., labor, material, energy, or cost & time together), optimizing a joint impurity at splits; in practice, many deployments train one boosted model per target and reconcile outputs post-hoc when correlations are strong [29].

Industrial evidence. Across manufacturing and related domains, ensembles consistently reduce cost/time estimation error versus classical regression, while retaining auditability via feature importance or SHAP-style attributions [23][24]. Reported use cases include component and total-cost prediction, cycle/lead-time estimation, and maintenance-duration/cost modeling on noisy, mixed-type datasets [17,27].

2.4 Artificial Neural Network regressions

Artificial Neural Networks (ANNs) learn distributed nonlinear mappings from features to targets via stacked affine–nonlinear layers. A two-hidden-layer MLP for regression can be written as:

$$\hat{y} = W_3\sigma W_2\sigma W_1 x + b_1 + b_2 + b_3, \tag{7}$$

With parameters W and b, nonlinearity σ, and training by backpropagation to minimize, e.g., mean-squared error $\frac{1}{n}\sum_i(y_i - \hat{y}(x_i))^2$. ANNs naturally extend to multi-output regression by using multiple neurons in the output layer (e.g., simultaneous prediction of cost and time, or cost components).

ANNs excel when relationships are highly nonlinear or inputs are high-dimensional (CAD/3D geometry, G-code/process traces, multi-sensor time series). CNNs can learn geometry-aware features directly from drawings/3D models for cost prediction; RNN/LSTM variants capture temporal shop-floor patterns for time-to-completion and downtime estimation [28]. With sufficient data and regularization, ANNs often match or exceed tree ensembles on complex cost/time tasks; hybrid stacks (e.g., boosting + shallow nets) are also reported. Interpretability can be improved with post-hoc explainers (SHAP/LIME) to support quotation and planning decisions.

Case studies show ANNs outperforming parametric cost models and classical regression for complex products and life-cycle cost estimation (e.g., electrical motors), and improving schedule/duration predictions where many features jointly drive outcomes. Practical constraints include data volume/quality and careful regularization; transfer learning or simpler architectures mitigate small-nnn scenarios.

On structured tabular manufacturing data, ensembles (RF, GBM/XGBoost) are strong defaults: they handle mixed features, capture interactions with modest tuning, and offer explainability often yielding the best accuracy among non-deep methods [23,24,25,27]. ANNs are preferable as data richness grows or inputs become high-dimensional/ unstructured (images/3D, sequences), and when predicting multiple correlated targets in one pass [17,26,28,29]. In practice, organizations often combine them: ensembles for robust baselines and governance; ANNs (or hybrids) to push accuracy where complexity warrants.

3 Case Studies and Approches

3.1 Cost Estimation

Predictive cost modelling is among the most frequent and business critical applications of regression methods in industrial management. It enables organisations to anticipate the cost profile of bids, design alternatives, or process changes and to price, budget, and seek savings with substantially less latency than ex-post costing or purely analogical expert judgement. With the diffusion of Industry 4.0 data infrastructures, the centre of gravity has shifted from retrospective calculation to proactive estimation: supervised machine learning (ML) models learn cost–driver relationships from historical enterprise data and update estimates as new evidence arrives. Recent reviews across project-driven and manufacturing contexts converge on this trajectory, highlighting ML/DL models' gains in accuracy, scalability, and the capacity to capture nonlinear interactions that traditional parametric models miss [17].

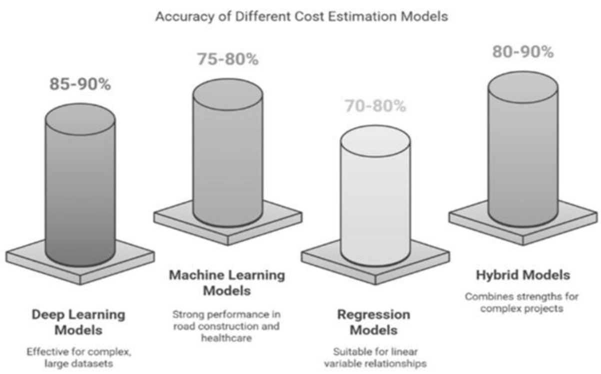

Fig. 2 *Accuracy of Different Cost Estimation Models (Shamim et al., 2025)*

In complex, high-mix environments, ensemble and deep models have been advanced to integrate signals from ERP/MES systems and shop-floor telemetry. Brown (2025) proposes an AI-driven framework that connects ERP-style features (materials, labour hours, machine time, overheads, batch size, design complexity) to total-cost prediction via Random Forests and sequence models, reporting reductions in estimation error (up to ~30%) and emphasising interpretability through feature-importance diagnostics. A representative implementation using a Random Forest Regressor on an ERP-like synthetic dataset attained a mean absolute error (MAE) of ≈ \$960 and $R^2 \approx 0.93$, while ranking labour hours and machine time as the dominant drivers insightful for cost engineering and margin protection. Beyond point accuracy, the work stresses integration concerns (data quality, feature selection, model transparency, and ERP/MES plumbing), positioning explainable, modular AI as a practical path for operations and finance teams [30].

For many SMEs, historical datasets are modest, noisy, and heterogeneous. Hennebold et al. (2022) address early-phase product development costing with a deliberately parsimonious feature set (basic concept parameters and firm-level indicators) and compare a spectrum of regressors – linear regression (LR), k-nearest neighbours (KNN), Random Forest (RF), and Support Vector Regression (SVR) on a compact dataset (81 finished products). The study argues for systematic feature definition, careful validation (MAE/RMSE), and an evidence-driven choice of algorithmic complexity, noting that under small-n conditions simple methods (e.g., KNN, LR) can be competitive with, or even preferable to, complex models that overfit. This "model parsimony first" principle resonates with broader ML guidance for industrial applications [17,18].

Cost estimation rarely concerns a single scalar; it often decomposes into interdependent components (e.g., labour, spares, downtime). He et al. (2023) demonstrate multi-target regression for scheduled warship maintenance, leveraging relationships among targets to improve simultaneous prediction of several cost-relevant quantities a structure that can reduce propagation errors when totals are aggregated from parts. Multi-Target Regression… As multi-output models share representation across targets, they are especially attractive when sub-costs are correlated and data per target are limited [29].

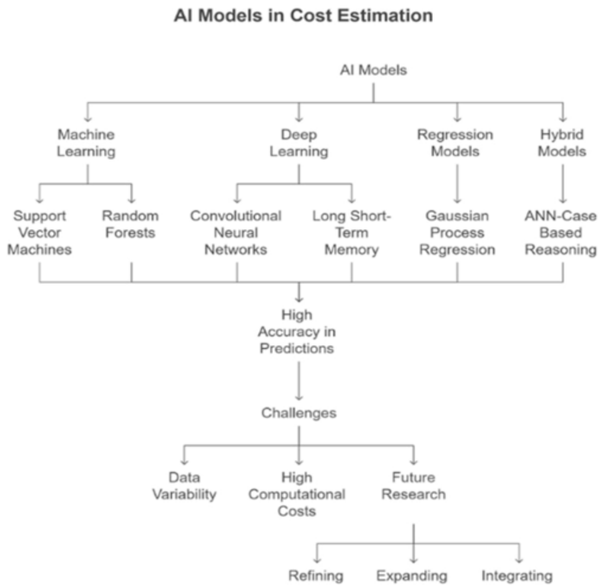

Fig. 3 *AI Models in Cost Estimation (Shamim et al., 2025)*

Where the economic asymmetry of under vs over estimation matters, classification style decision layers can be made explicitly cost-aware. In a batch process case (thermocouple alloy production), engineeris combine batch-data analytics with cost-sensitive learning: predicted compliance decisions are filtered by a probability threshold to trade off accuracy and coverage, and a cost matrix converts misclassification patterns into monetary "relative cost." A Random Forest-based policy with a 65% confidence threshold achieved a *relative cost* of 26% versus the baseline, providing a tangible, economics-driven operating point instead of purely accuracy-driven tuning. The formalism confusion and cost matrices and a scalarised cost functional clarifies how to select cut-offs for the plant's objectives and to handle "out-of-coverage" cases (accept vs reject) [31].

Beyond tabular ERP attributes, cost models increasingly ingest modality-rich inputs. Text or code-based specifications (e.g., customer descriptions, CAD/CAM artefacts, or G-code) can seed automated feature extraction, enabling rapid quotation for customised products via NLP and ML pipelines an emerging line exemplified by work on cost estimation from free-form customer inputs. This aligns with the broader literature that argues for hybrid pipelines: rules or parametrics where physics is strong, ML where patterns are high-dimensional and nonlinear, and explicit uncertainty quantification where decisions carry asymmetric financial risk.

Practically, organisations progress along a deployment ladder. At one end are spreadsheet-resident regressors with routine data refresh; at the other are services connected to ERP/MES with scheduled retraining, monitoring, and explainability dashboards. Brown (2025) explicitly discusses this integration continuum and the cultural/technical prerequisites for adoption (data governance, interpretability, SME readiness). In parallel, cost-sensitive decision-support exemplars show how to translate model confidence into economically meaningful actions on the shop floor [30,31].

The case studies above map a design space rather than a single "best" model. In data-rich, heterogeneous settings, ensembles and deep learners deliver strong accuracy and actionable driver insights; in low-data regimes, simpler regressors with deliberate feature engineering and robust validation are often preferable. When decisions are binary (accept/reject, bid/no-bid) and errors are financially asymmetric, cost-sensitive learning converts predictive scores into business-optimal policies. Multi-target structures help when costs decompose into correlated parts. Across settings, the consistent pattern is to match model complexity to data realism, couple accuracy with economics, and invest early in feature pipelines and interpretability to ensure adoption by engineering and finance stakeholders.

3.2 Time Estimation

Predicting the duration of operations and the remaining completion times of orders is one of the key tasks of cross-industry management, especially in make-to-order mode. A more accurate estimate of "when it will be done" allows for better capacity planning, guaranteeing deadlines and reducing WIP and downtime. However, traditional approaches (norms, deterministic calculations in planning/simulation tools) have difficulty absorbing the dynamics of real operations changes in priorities, process variability, sudden failures or urgent insertions. Therefore, data-driven machine learning (ML)-based prediction, often linked to real-time data (MES/IIoT), is gaining ground. Reviews in recent years confirm that ML approaches are penetrating production planning and control in addition to maintenance and quality, with the use of deep sequential models and tree ensembles in time estimation and scheduling tasks growing significantly.

A powerful example is the MORCT (Multiple Orders Remaining Completion Time) problem in discrete shops, where the remaining times for multiple work-in-progress orders competing for machine capacity need to be predicted simultaneously. Liu et al. propose a framework, DMTR-LSA (Deep Multi-Target Regression with LSTM Self-Attention), which dynamically composes a dataset from multiple production sources reflecting the current shop floor status (machine load, WIP, etc.) and jointly predicts the completion times of all orders. In real operation, the model reduced the average prediction error by ≈ 8.98% compared to established methods by combining LSTM (sequential dependency) with self-attenuation (weighting of important sections). It also exhibits adaptivity for faults or urgent insertions that traditional (statistical or simulation) approaches often fail to capture without recalibration [4,33].

Conceptually, this approach uses multi-target/multi-output learning: the model predicts a vector of target variables (times for multiple jobs) at once, and can thus exploit cross-correlation and shared representation. Theoretical work shows that joint target modeling leads to better accuracy in many domains than a set of independent single-target models (chains/ensemble variants) [34][38]. In addition to sequential networks, gradient-boosting methods augmented with mechanisms for sharing information between outputs can be considered for multi-output regression, achieving competitive results with lower data and infrastructure requirements especially useful in SMEs [33].

From a "what works in practice" perspective, systematic reviews confirm that in industrial prediction tasks, tree ensemble models (RF, XGBoost, GB) are among the most commonly and successfully deployed; sequential DL (LSTM/CNN) is dominant where rich temporal sensing is available. The choice of architecture is contextual: environments with fewer sensors tend to have more robust ensembles, data-rich environments benefit from more complex DL models. For evaluation, it is also recommended to abandon pure classification metrics in favor of regression error (MAE/RMSE) and to consistently report data (im)balance, as aggregate accuracy can be misleading in the presence of class imbalance [35,37].

For SMEs, a full-fledged implementation of DMTR-LSA (streaming, inference orchestrator, MLOps) is often unrealistic. A practical way is the stacked single-target approach: for key operations, train lightweight regressors (e.g., tree/GB) on order and machine parameters (material, batch, tool, production shift), aggregate their outputs, and incorporate uncertainty estimation into the plans. This route usually yields a leap in refinement over normative models, and is also compatible with Excel/ERP. Empirical work from the Czech environment confirms that even relatively simple models (RF, GB, SVR) with well-designed features can significantly improve the prediction of operations and lead times in shop floor management [33].

In larger enterprises, lead time prediction is integrated into MES/APS: sensor data (cycle times, status signals, OEE) flow on-line into models that continuously update job completion estimates, visualize them to dispatchers and pass them to dispatchers/reschedulers. The trend comes from the PdM field, where RUL is commonly formulated as a regression problem and sequential architectures (LSTM/CNN) excel at capturing degradation patterns; the same logic (sequence → regression time to event) is transferable to production time prediction [35,36]. Moreover, the new SLR documents the shift from lab to real deployments: a significant fraction of the studies are running in "live" operation, connected to maintenance/work-order systems and reporting tangible benefits (reduced downtime, higher OEE).

4 Conclusion

This paper has shown that machine learning regression techniques provide a coherent framework for two key managerial variables in manufacturing: cost and time. For cost, we highlighted the shift from expert spreadsheets and rules to data-driven models (from classical regressions to trees to neural networks) that can capture non-linearities and interactions of input features and improve estimation accuracy at the operation, order and portfolio level. For time, we have shown that the same regression principles (including advanced sequential architectures) allow us to predict operation durations, lead times, and remaining completion times in a multi-step pro-environment. Specific industry studies point to significant improvements over traditional approaches when models are linked to operational data and can handle non-stationarities and operational "noise".

Of particular note is the multi-target/multi-output approach, which simultaneously predicts multiple outputs (e.g. sub-cost items or remaining times of multiple jobs) and exploits their interrelationships. Work in MTR shows that sharing representations across targets and explicitly modeling correlations (including suppressing structure-duct noise) leads to more robust and accurate prediction than separate single-target models even in small-sample environments. These findings also frame industrial applications where the overall target (e.g., total cost or total time) is obtained by summing more accurately predicted sub-targets.

For production time prediction, deep sequential models (LSTMs with self-attenuation) offer a way to capture shop floor dynamics and resource competition between parallel orders. A MORCT study for discrete manufacturing shows that real-time multi-objective prediction of remaining times is feasible and reduces error by units to tens of percent compared to traditional methods; this translates directly into quality of scheduling and meeting deadlines. And multi-target regressions... At the same time, databases from real companies confirm that even more "classical" regression models suitably adjusted for outliers, supplemented with textual features and uncertainty estimation can significantly refine times at the operations level in data-constrained environments.

Another natural link is predictive maintenance. Although it primarily targets RUL/time-to-failure, its outputs can be viewed as explanatory variables in regression models for machine availability and hence for cost and time predictions. Recent reviews confirm the trend towards time series models (LSTM, attention) and towards combinations with physical knowledge, which increases robustness under changing conditions.

In terms of practical implementation, scalability is key: large enterprises can integrate II-oT/MES/APS into a single chain and update predictions online; SMEs can start with simpler models (e.g., trees/GBM in Excel or ERP) and gradually add features and data sources. The literature on cost and time suggests that "staged" adoption from single-target regressions to multi-target and online learning yields tangible gains in the early stages.

Limitations of current approaches include the demands on data quality (missingness, error rates, shifts in distributions), the explainability of complex models, and the cost of integration into planning processes. Therefore, it makes sense to Combine ML with its simulation/digital twin: simulation generates scenarios and constraints, ML provides quick estimates and corrections based on reality. Moreover, structurally rich data (panels, sequences, images) call for me-methods that work naturally with tensor representations and can regularise over dimensions (time × machine × variance), supporting modern approaches to tensor regression.

In conclusion, we see the greatest potential in unifying cost and time predictions into a single multi-objective framework that shares representations across sub-items and operations and better captures real operational linkages (queues, shifts, capacity). This approach is intended to build on hybrid, physically-informed models combining queueing simulation/theory with regression (including tensor or kernel formulations for high-dimensional, structured data), with on-line learning with "concept drift" detection and management and integration of PdM outputs (RUL) into time predictions and risk reserves as standard. The key for planning is to move to probabilistic, well-calibrated outputs (intervals, error distributions), while strengthening explainability (SHAP, contrastive explanation) and "human-in-the-loop" active learning. Across disciplines, we then do-recommend standardizing data models and building open benchmarks for cost and time tasks (including sequential and multi-target scenarios). This will create a continuous pipeline that combines data, simulation, and learning into a single decision mechanism and takes the accuracy, robustness, and practical applicability of predictions another order of magnitude further.

This text is a pilot synthesizing study and forms the first part of a planned series. In the follow-up applied research, the procedures and ideas formulated here will be tested on real production data and implemented in transparent regression pipelines; emphasis will be placed on replicable methodology, generalization metrics (train-validation-test) and case studies, especially in the environment of small and medium-sized enterprises. The results will be published in the future sequels.

References

[1] BODENDORF, F., FRANKE, J. (2021). A machine learning approach to estimate product costs in the early product design phase: a use case from the automotive industry. In: *Procedia CIRP*, Vol. 100, pp. 643 – 648. Elsevier B. V. Netherlands. ISSN 2212-8271.

[2] IŞIK, K., ALPTEKIN, E., S. (2022). A benchmark comparison of Gaussian process regression, support vector machines, and ANFIS for man-hour prediction in power transformers manufacturing. In: *Procedia Computer Science*, Vol. 207, pp. 2567 – 2577. Elsevier B. Netherlands. ISSN 1877-0509.

[3] BODENDORF, F., MERKL, P., FRANKE, J. (2021). Intelligent cost estimation by machine learning in supply management: A structured literature review. In: *Computers and Industrial Engineering*, Vol. 160. Elsevier Ltd. United Kingdom. ISSN 0360-8352.

[4] LIU, B., JIANG, Z., H. (2005). The man-hour estimation models & its comparison of interim products assembly for shipbuilding. In: *International Journal of Operations Research*, Vol. 2, No. 1, pp. 9 – 14. Inderscience Enterprises. United Kingdom. ISSN 1745-7645.

[5] HUR, M., LEE, S., K., BONGSEOK, K., CHO, S., LEE, Do., LEE, Da. (2015). A study on the man-hour prediction system for shipbuilding. In: *Journal of Intelligent Manufacturing*, Vol. 26, No. 6, pp. 1267 – 1279. Springer. Netherlands. ISSN 0956-5515.

[6] YU, T., CAI, H. (2015). The Prediction of the Man-Hour in Aircraft Assembly Based on Support Vector Machine Particle Swarm Optimization. In: *Journal of Aerospace Technology and Management*, Vol. 7, No. 1, pp. 19 – 30. Departamento de Ciencia e Tecnologia Aeroespacial. Brazil. ISSN 1984-9648.

[7] WEN, J., LI, S., LIN, Z., HU, Y., HUANG, CH. (2012). Systematic literature review of machine learning based software development effort estimation models. In: *Information and Software Technology*, Vol. 54, No. 1, pp. 41 – 59. Elsevier. Netherlands. ISSN 0950-5849.

[8] QIN, J., HU, F., LIU, Y., WITHERELL, P., WANG, C., C., L., ROSEN, D., W., SIMPSON, T., W., LU, Y., TANG, Q. (2022). Research and application of machine learning for additive manufacturing. In: *Additive Manufacturing*, Vol. 52. Elsevier. Netherlands. ISSN 2214-8604.

[9] CHAN, S., L., LU, Y., WANG, Y. (2018). Data-driven cost estimation for additive manufacturing in cybermanufacturing. In: *Journal of Manufacturing Systems*, Vol. 46, pp. 115 – 126. Elsevier. Netherlands. ISSN 0278-6125.

[10] DENG, S., YEH, T., S. (2011). Using least squares support vector machines for the airframe structures manufacturing cost estimation. In: *International Journal of Production Economics*, Vol. 131, No. 2, pp. 701 – 708. Elsevier. Netherlands. ISSN 0925-5273.

[11] SAJADFAR, N., MA, Y. (2015). A hybrid cost estimation framework based on feature-oriented data mining approach. In: *Advanced Engineering Informatics*, Vol. 29, No. 3, pp. 633 – 647. Elsevier. United Kingdom. ISSN 1474-0346.

[12] GUNDUZ, M., UGUR, L., O., OZTURK, E. (2011). Parametric cost estimation system for light rail transit and metro trackworks. In: *Expert Systems with Applications*, Vol. 38, No. 3, pp. 2873 – 2877. Elsevier. United Kingdom. ISSN 0957-4174.

[13] HAMRANI, A., AGARWAL, A., ALLOUHI, A., MCDANIEL, D. (2024). Applying machine learning to wire arc additive manufacturing: a systematic data-driven literature review. In: *Journal of Intelligent Manufacturing*, Vol. 35, No. 6, pp. 2407 – 2439. Springer. Netherlands. ISSN 0956-5515.

[14] NGUYEN, T., H., V., HUANG, P., M., CHIEN, CH., F., CHANG, CH., K. (2023). Digital transformation for cost estimation system via meta-learning and an empirical study in aerospace industry. In: *Computers and Industrial Engineering*, Vol. 184. Elsevier. United Kingdom. ISSN 0360-8352.

[15] NING, F., SHI, Y., CAI, M., XU, W., ZHANG, X. (2020). Manufacturing cost estimation based on the machining process and deep-learning method. In: *Journal of Manufacturing Systems*, Vol. 56, pp. 11 – 12. Elsevier. Netherlands. ISSN 0278-6125.

[16] ELMOUSALAMI, H., H. (2020). Artificial Intelligence and Parametric Construction Cost Estimate Modeling: State-of-the-Art Review. In: *Journal of Construction Engineering and Management*, Vol. 146, No. 1. American Society of Civil Engineers. United States. ISSN 0733-9364.

[17] SHAMIM, M., M., I., HAMID, A., B., A., NYAMASVISVA, T., E., RAFI, N., S., B. (2025). Advancement of Artificial Intelligence in Cost Estimation for Project Management Success: A Systematic Review of Machine Learning, Deep Learning, Regression, and Hybrid Models. Online. *Modelling*. Vol. 2025, no. 6, article 35. ISSN 2673-3951. https://doi.org/https://doi.org/10.3390/modelling6020035.

[18] HENNEBOLD, CH., KLÖPFER, K., LETTENBAUER, P., HUBER, M. (2022) Machine Learning based Cost Prediction for Product Development in Mechanical Engineering. Online. *Procedia CIRP (55th CIRP Conference on Manufacturing Systems)*. Article 107, s. 264-269. ISSN 2212-8271. https://doi.org/10.1016/j.procir.2022.04.043.

[19] LIU, J., ZHU, C., LONG, Z., LIU, Y. (2021). *"Tensor Regression", Foundations and Trends® in Machine Learning*. 14. No. 4, pp 379-565. Hanover, MA 02339: now Publishers. ISBN 978-1-68083-887-9. http://dx.doi.org/10.1561/2200000087.

[20] LIU, M., L., ZHANG, J., QIN, S., ZHANG, K., WANG, S., et al. (2025). A multi-target regression-based method for multiple orders remaining completion time prediction in discrete manufacturing workshops. Online.

Measurement. Vol. 2025, no. 242, article 116231. ISSN 0263-2241. https://doi.org/https://doi.org/10.1016/j.measurement.2024.116231.

[21] NIAZI, A., DAI, J., S., BALABANI, S., SENEVIRATNE, L., (2021). "Product Cost Estimation: Technique Classification and Methodology Review." *ASME*. J. Manuf. Sci. Eng. May 2006; 128(2): 563–575. https://doi.org/10.1115/1.2137750

[22] BREIMAN, L., FRIEDMAN, J., OLSHEN, R., A., STONE, C., J. (1984). *Classification and Regression Trees* (1st ed.). Chapman and Hall/CRC. https://doi.org/10.1201/9781315139470

[23] MONDON, C. (1984). Classification and regression trees.

[24] BREIMAN, L. *Random Forests*. (2001). Machine Learning 45, 5–32 https://doi.org/10.1023/A:1010933404324

[25] FRIEDMAN, J. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics,* 29, 1189-1232. http://dx.doi.org/10.1214/aos/1013203451

[26] TIANGI, CH., GUESTRIN, C. (2016). XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16). Association for Computing Machinery, New York, NY, USA, 785–794. https://doi.org/10.1145/2939672.2939785

[27] LESZCZYŃSKI, Z., JASIŃSKI, T. (2020). Comparison of Product Life Cycle Cost Estimating Models Based on Neural Networks and Parametric Techniques—A Case Study for Induction Motors. Sustainability. 12. 8353. 10.3390/su12208353.

[28] KO, H., S., SEO, H., G. (2025). A Machine Learning Approach for Manufacturing Cost Estimation. Online. Roč. 2024, s. 181-214. ISSN 1598-978X. From: https://doi.org/https://doi.org/10.31507/KJMAR.2024.12.24.3.181. [cit. 2025-07-26].

[29] NING, F., QU, H., SHI, Y., CAI, M., Xu, W. (2022). Feature-Based and Process-Based Manufacturing Cost Estimation. Machines. 10. 319. 10.3390/machines10050319.

[30] HE, D., SUN, S., XIE, L. (2023). Multi-target regression based on multi-layer sparse structure and its application in warships scheduled maintenance cost prediction. *Applied Sciences, 13*(1), 435.

[31] BROWN, B. (2025). AI-Driven Cost Estimation Models for Complex Manufacturing Systems. *Online*. Roč. 2025. From: https://www.researchgate.net/publication/393661904.

[32] MÄHLKVIST, S., EJENSTAM, J., KYPRIANIDIS, K. (2023). Cost-Sensitive Decision Support for Industrial Batch Processes. Online. *Sensors*. Vol. 2023, no. 23, article 9464. ISSN 1424-8220. From: https://doi.org/https://doi.org/10.3390/s23239464.

[33] LIU, M., L., ZHANG, J., QIN, S., ZHANG, K., WANG, S., et al. (2025). A multi-target regression-based method for multiple orders remaining completion time prediction in discrete manufacturing workshops. Online. *Measurement*. Vol. 2025, no. 242, article 116231. ISSN 0263-2241. From: https://doi.org/https://doi.org/10.1016/j.measurement.2024.116231.

[34] ANTONENKO, E. (2023). Multi-target Learning and Prediction : Novel Methods and Application-s. Artificial Intelligence. [cs.AI]. Institut Polytechnique de Paris. ISBN NNT : 2023IPPAX098.

[35] TSALLIS, CH., PAPAGEORGAS, P., PIROMALIS, D., MUNTEANU, R., A. (2025). Application-Wise Review of Machine Learning-Based Predictive Maintenance: Trends, Challenges, and Future Directions. Online. *Applied Sciences*. 2025, No. 15, article 4898. ISSN 2076-3417. Available: https://doi.org/https://doi.org/10.3390/app15094898.

[36] ZHANG, M., SUN, Z., ZHENG, W. (2025). Fault Diagnosis of Electric Motor Rotor System Based on Feature Exreaction & CNN-BiGRU-Attention. *Manufacturing Technology*. Vol. 25, No. 4. ISSN 1213–2489. DOI: 10.21062/mft.2025.048

[37] CÉZOVÁ, E. (2025). Demonstration of Neural Network in Prediction of Bearing Lifetime. *Manufacturing Technology*. Vol. 25, No. 2. ISSN 1213–2489. DOI: 10.21062/mft.2025.017

[38] LUSI, N., FIVERIATI, A., AFANDI, A., WEDARMA, G., N., B., C., S., YULIANDOKO, H., DARSIN, M., QUTABA, S. (2023). Predictive Modelling On Machining Performance of ECDM Using Artificial Neural Network and Particle Swarm Optimization. *Manufacturing Technology*. Vol. 23, No. 5. ISSN 1213–2489. DOI: 10.21062/mft.2023.076