Technologies for Manufacturing as a Service Ecosystems

Deliverable D3.3

D3.3. Analytics for resource-subservice matching and service composition

WP3: DT Modelling, Operation and Governance for resilient value networks

Editor:	Konstantinos Kaparis
Lead beneficiary:	UoM
Version:	1.0
Status:	Final
Delivery date:	30/06/2025
Dissemination level:	PU (Public)

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101138517.

Deliverable Factsheet

Grant Agreement No.	101138517
Project Acronym	Tec4MaaSEs
Project Title	Technologies for Manufacturing as a Service Ecosystems
Start date	01/01/2024
Duration	36 months

Deliverable Name	D3.3 Analytics for resource-subservice matching and service composition
Related WP	WP3: DT Modelling, Operation and Governance for resilient value networks
Due Date	30/06/2025

Author(s)	Konstantinos Kaparis (UoM)
Contributor(s)	Kyriakos Bitsis, Andreas Georgiou, Dimitris Thomaidis, Georgios Tsaples, Angelo Sifaleras, Ioannis Konstantaras (UoM), Patricia Casla, Alvaro Garcia, Itxaso Cascon, Jon Ander Sarasua (TEK), Sotiria Antaranian, Giorgos Triantafyllou (ATC), Metin Turkay, Mert Can (SMO), Ioannis Mourtos, Ioannis Augerinos (AUEB), Kostas Kalaboukas, Dimitris Verlekis (MAG)
Reviewer(s)	Eleni Gkinou (MAG)
Approved by	All partners

Document History

Version	Date	Author(s)	Organization	Description
0.1	08/05/2025	Konstantinos Kaparis	UoM	ToC
0.2	12/06/2025	Patricia Casla, Alvaro Garcia, Itxaso Cascon, Jon Ander Sarasua	TEK	Contribution to service decomposition, Bill of Process Generator: requirements extraction and process decomposition and processing time estimation.
0.3	19/06/2025	Mert Can Özcan, Metin Türkay	SMO	Analytics service (forecasting) of VN1
0.4	20/06/2025	Sotiria Antaranian	ATC	Updated 2.1.2.2
0.5	25/06/2025	Konstantinos Kaparis, Kyriakos Bitsis, Andreas Georgiou, Dimitris Thomaidis, Georgios Tsaples	UoM	Complete draft for internal review
0.6	27/06/2025	Eleni Gkinou	MAG	Peer review
0.7	30/06/2025	Armend Duzha	MAG	Quality check
1.0	30/06/2025	Konstantinos Kaparis	UoM	Final version ready for submission

Executive Summary

This deliverable presents the initial version of analytics services designed to support resource–subservice matching and service composition within the Tec4MaaSEs architecture. These services enable the flexible configuration of MaaS value chains by processing consumer requests and identifying appropriate combinations of providers and resources across the platform's value networks. The analytics engine functions as a modular component, interfacing with the data space, consumer request specifications, and the platform's optimization service.

At its core, the analytics layer delivers provider-specific insights, system-level information, and post-processing outputs that assist in capability identification, task decomposition, resource matching, and the composition of services. The resulting outputs are structured as inputs to the optimization process, linking analytical results with real-time decision-making support.

This first version focuses on implementations in VN1 and VN2, where analytics services are tailored to the specific planning level, volume–variety dynamics, and operational requirements of each pilot. For each service, the deliverable outlines functional specifications, pilot-specific use cases, and data exchange mechanisms. A dedicated section also describes how analytics outputs interact with the optEngine and contribute to the configuration of production scheduling workflows.

Overall, the analytics services presented here establish a foundation for contextualized, explainable, and efficient decision-making in MaaS environments. These services will be further extended, validated, and refined in D3.4.

Table of Contents

List of Figures

List of Tables

Acronyms and Abbreviations

Acronym	Description
AM	Additive Manufacturing
BoPG	Bill of Process Generator
CAD	Computer-Aided Design
EB	Electronic Boards
EB	Electronic Boards
QoS	Quality of Service
VN	Value Network

1 Introduction

1.1 Purpose and Scope

This deliverable includes the initial design and implementation of analytics services that supports the process of resource-subservice matching and service composition within the Tec4MaaSEs platform. These analytics services constitute a critical functional layer that enables data-driven decision-making across the platform's value networks. Their main objective is to transform data from consumer requests, provider capabilities into actionable structures and artifacts, such as decomposed service modules, matched resources and ranked value chains.

The scope of D3.3 spans descriptive, predictive, and post-processing analytics, supporting capability detection, demand forecasting, decomposition, and multi-criteria filtering and ranking of service compositions. The analytics layer is modular and pilot-specific, reflecting the varying planning levels, operational complexities, and customization needs of the three VNs.

This deliverable is directly related to:

- D2.1 (Requirements and functional design): from which user stories, interaction diagrams, and planning-level distinctions per value network are derived.
- D2.2 (Platform architecture and information models): which provides the technical integration framework for the analytics services.
- D3.2 (Design and implementation of optimization services): with which D3.3 maintains a close functional and data-level interface, particularly for exchanging inputs/outputs with the Optimization Engine.
- D3.4 (Final integrated decision-support analytics services): where the analytics modules introduced in D3.3 will be further extended, validated, and refined based on full-scale testing.

The work in D3.3 also contributes to WP5 (Integration and Demonstration), providing pilot-specific analytics support for evaluation against KPIs, as well as WP4 (Digital Twin services), by taking advantage of data arising DT instances for matching and predictive modeling tasks.

1.2 Structure of the Document

The structure of this deliverable reflects the functional layering and pilot-specific orientation of the analytics services developed in WP3. It is organized as follows: Section 2 introduces the positioning of analytics within the overall platform architecture, highlighting their descriptive, predictive, and post-optimization components. It explains how these services support resource-subservice matching and service composition, and how they interact with the optimization engine. Topics such as trust, data flow, and modularity are also addressed. Section 0 presents analytics-related insights and developments for VN1 and VN2, including their context, relevant user stories, key performance indicators (KPIs), and planning-level characteristics. It also outlines how analytics services are customized to match the operational needs of each pilot. Section 0 summarizes the main contributions of the deliverable and outlines next steps toward D3.4, where the analytics layer will be fully integrated and validated within the platform environment. The report concludes with a list of references.

2 Analytics Services

2.1 Role of Analytics in the T4M architecture

In a Manufacturing-as-a-Service (MaaS) ecosystem, analytic services form the intelligence core that drives data-informed decision-making, operational efficiency, and adaptability. MaaS platforms connect resource providers and service consumers through a cloud-based infrastructure, where analytics serve as the enabler of real-time insights and strategic optimization. Such services should be able to process vast streams of data from machinery, production lines, supply chains, and consumer interactions to inform key areas such as performance monitoring, predictive maintenance, quality control, and demand forecasting. Within the Tec4MaaSEs platform, analytics take on an even more foundational role. They empower dynamic decisions that are essential for service decomposition, matching, and composition, ensuring that resources are allocated and configured with maximum efficiency. Designed to interact with Digital Twin representations of manufacturing resources, these analytic services provide explainable, context-aware recommendations that guide and support the platform's optimization engine. This integration enables the formation of optimal, flexible manufacturing configurations, tailored to evolving demands and constraints. By transforming raw data into actionable intelligence, the analytic layer in Tec4MaaSEs not only enhances operational performance but also supports the creation of smarter, more resilient, and innovation-driven manufacturing ecosystems.

To realize the vision of next-generation MaaS platforms, Tec4MaaSEs promotes a shift from rigid production models to a more dynamic, data-driven approach. This approach is guided by the unique characteristics of each domain registered on the platform, leveraging domain-specific volume–variety profiles to deliver tailored analytics that align with the capabilities and needs of individual entities.

The platform's design is validated through three pilot cases; each selected to reflect this paradigm and demonstrate scalability across diverse value networks. These pilots span distinct positions within the volume–variety space and correspond to different quadrants of the Product-Process Matrix (PPM) as defined by Hayes and Wheelwright (1979). Their diversity highlights that a one-size-fits-all analytic solution is not viable in a MaaS marketplace. Instead, Tec4MaaSEs implements a foundational analytics architecture capable of supporting domain-specific analytic services tailored to contextual needs. These services vary

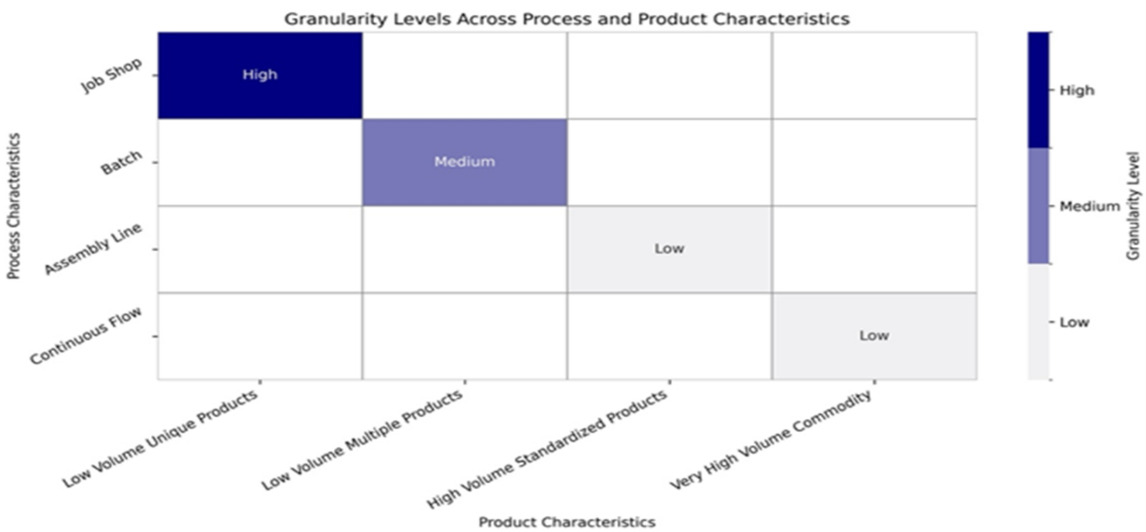

Figure 1: Volume vs Variety / Product vs Process granularity levels

significantly, from supporting high-volume, low-variability operations to enabling highly customized, low-volume production, underscoring the importance of context-aware, modular analytics (see Figure 1).

The three Value Networks (VN1–VN3) implemented in Tec4MaaSEs exemplify distinct levels of planning in operations management—operational, aggregate, and strategic—while illustrating a wide range of volume–variety dynamics and process granularities:

- **VN1 (Operational Planning):** Focuses on day-to-day coordination of Electronic Board (EB) production and distribution among four stakeholders. Operating in a high-volume, low-variety environment, the network uses analytics to optimize resource allocation, balance supply and demand, and improve responsiveness in a complex supply chain.
- **VN2 (Aggregate Planning):** Involves three stakeholders engaged in additive and subtractive manufacturing for mold production. Operating in a low-volume, high-variety context, this network relies on analytics for dynamic service decomposition, task allocation, and capacity matching—minimizing idle time and enabling on-demand workflow reconfiguration.
- **VN3 (Strategic Planning):** Led by a contractor managing large-scale energy infrastructure projects, this network focuses not on production but on transforming procurement and negotiation workflows. The platform supports standardized, digital exchange of rich technical information with multiple equipment suppliers, replacing paper-based processes and enhancing traceability and coordination.

Together, these pilots demonstrate how Tec4MaaSEs supports diverse planning and decision-making needs across heterogeneous manufacturing contexts. They reinforce the necessity of a modular, context-sensitive analytics layer that can adapt to varying complexities, process granularities, and value network structures within the broader MaaS ecosystem. The analytics services are structured around three core pillars, tailored to the specific requirements of each Value Network. Nevertheless, as will be discussed later, although VN1 and VN2 exhibit different needs, these services show limited value-added potential in the context of VN3.

1. **Analytics for Composition Pre-processing:** This component focuses on analytics that support the preparatory stages leading to the derivation of optimal compositions (Section 2.1.2). It includes tasks such as service decomposition, provider matching, demand forecasting, and dimensionality reduction. These activities streamline and enhance the subsequent optimization process by reducing complexity and aligning inputs with the system's dynamic requirements (applicable to VN1 and VN2).
2. **Analytics for Composition Post-Processing:** The third pillar addresses analytics applied after the optimization phase (Section 2.1.3) to evaluate, interpret, and refine its outcomes. These post-processing methods further simplify the solution space by handling constraints and conditions in a computationally efficient manner, thereby avoiding unnecessary increases in optimization complexity. This step ensures that results are actionable, robust, and aligned with real-world implementation needs (applicable to VN1 and VN2).
3. **Descriptive Analytics and Visualization:** It provides foundational insights by visualizing system performance metrics, key operational indicators and general information. This enables stakeholders to obtain a comprehensive understanding of the current state of the system, serving as a basis for informed decision-making (applicable for all VNs through platform horizontal functionalities). See Section 2.1.4 below.

By integrating these pillars, the analytics layer goes beyond supporting the optimization phase, as it fully embeds and structures it within a continuous, data-driven loop. Each pilot applies a tailored combination of analytics methods, aligned with its specific planning horizon, volume and variety characteristics, and operational complexity. This modular design ensures that analytics remain context-aware, contributing meaningfully to both system-level intelligence and the adaptability of the platform.

The subsection 2.1.1 briefly introduces the composition/optimisation service, as the first two pillars are closely intertwined with it. Therefore, we begin with a short overview of this service.

2.1.1 Composition service

The composition component in the Tec4MaaSEs platform is a central element within the context-driven, optimized (re-)configuration services. It plays a critical role in converting manufacturing service requests into production schedules by dynamically composing services across multiple providers in a value network. This functionality supports the broader goals of Manufacturing-as-a-Service (MaaS) by enabling flexible, efficient, and context aware production planning. The composition process begins with the receipt of structured service requests submitted by consumers. These requests typically include detailed product requirements, delivery constraints, and business objectives. The analytics module processes these inputs, filling in missing values and structuring the data into a standardized format that includes information such as operation sequences, machine eligibility, time windows, and preferences like locality or makespan (time that the last job is finished). Once structured, this data is passed in JSON format to the optimization engine for further processing. At the core of the composition component lies a set of mathematical models, primarily Constraint Programming (CP) and Mixed-Integer Linear Programming (MILP). These models are tailored to address distributed versions of classic flowshop and hybrid flexible flowshop scheduling problems. They incorporate a variety of constraints, such as machine eligibility, job precedence, provider availability, and transportation times. The models are designed to optimize multiple objectives simultaneously, including the minimization of makespan, earliness and tardiness, transportation-based locality weights, and counterbalance the job distribution across providers.

A key feature of the composition service is its ability to generate multiple alternative schedules in response to a single request. These alternative compositions support post-optimization processes where stakeholders can evaluate trade-offs between different objectives and collectively decide on the most suitable production plan. This is achieved by solving the scheduling model multiple times with varied objective weightings and incorporating constraints that prevent repetition of previous solutions. The approach ensures that the alternatives are meaningfully different, enabling more informed decision-making. Once stakeholders agree on a preferred schedule, the composition component facilitates its execution by communicating the plan to the relevant providers, potentially through their digital twins. The system supports ongoing feedback and real-time updates. If disruptions occur, such as changes in machine availability, the optimization cycle can be re-initiated to produce updated compositions that reflect the new context. This feedback loop enhances resilience and adaptability in dynamic manufacturing environments. Within the Tec4MaaSEs platform, the composition service is integrated with the Optimization Engine, a modular shell responsible for handling optimization job submissions, managing asynchronous execution, and storing solutions in the platform's data infrastructure. This architectural approach ensures scalability, robustness, and seamless interaction with other platform components, such as analytics and data registries.

The application of the composition component varies across the project’s three value networks. In VN2, which includes additive manufacturing and plastic injection moulding, the composition service is most extensively developed and tested, with complex models addressing multi-objective trade-offs. VN1, which focuses on electronic board production, shares a similar modelling approach but presents scalability challenges due to higher job volumes, which are mitigated through preprocessing techniques like job batching. In contrast, VN3 is centered on long-term procurement and strategic negotiations, so due to the customised nature of interactions makes the composition service less relevant.

2.1.2 Analytics for composition pre-processing

2.1.2.1 Task decomposition service

In the Tec4MaaSEs framework, decomposition analytics serve as a foundational mechanism for translating complex, heterogeneous manufacturing requests into modular subservices that can be dynamically matched and orchestrated. It is important to note that the decomposition phase is not always imperative in a MaaS framework. Its necessity depends on the characteristics of each manufacturing domain and our three VNs clearly demonstrate this distinction. The tool that is used within T4M for the automatic decomposition of comsumer's requests is the "Bill of Process Generator" and is applied in VN2. The use of Additive Manufacturing (AM) is becoming increasingly common in the production of certain mould components, particularly cavity inserts and metal structural parts. AM is valued for its ability to produce highly complex geometries, offer easy customization to specific user requirements, and provide flexibility and cost-efficiency. However, when producing metal parts, AM alone is often not sufficient to meet the stringent geometrical and surface quality requirements demanded by customers. As a result, additional finishing processes such as machining are typically required.

Machining itself is a complex process involving the precise removal of material to achieve the desired shape and surface finish. To manage this complexity, the machining workflow is usually decomposed into a series of well-defined operations, each serving a specific purpose. These operations commonly include milling, drilling, threading, grinding, and others, depending on the part's geometry and functional specifications. This process decomposition is critical for effective process planning, accurate toolpath generation, and appropriate machine and tool selection.

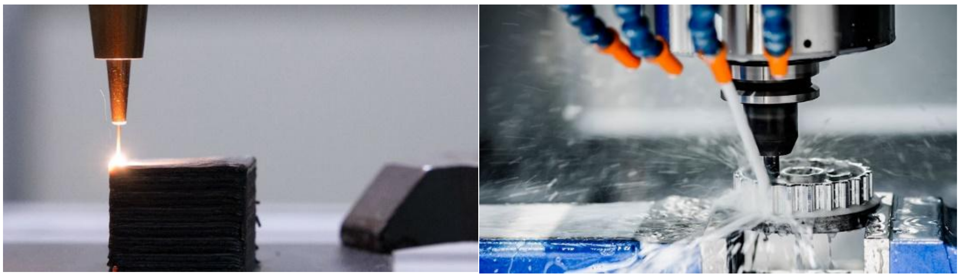

Figure 2: Examples of laser cladding (additive manufacturing) and milling (subtractive manufacturing)

One of the main challenges manufacturing companies face daily is decomposing the manufacturing process of a component into a sequence of operations and identifying suitable machines based on the technical requirements of the part. Currently, this remains a problem solved manually through the experience of skilled operators. Generally, the manufacturing company receives a blueprint of a part (3D CAD file). Based on the quality requirements and geometric characteristics of the part, the production team determines the sequence of processes (forging, milling, drilling, polishing, electrical discharge machining, etc.), considering both intrinsic limitations of each process and those derived from their own equipment. Subconsciously, the operator follows this sequence of actions:

1. Observes the geometry of the part: whether rotational, prismatic, or with hard-to-reach corners, and intuitively sketches out the main processes that should be involved. For instance, a rotational geometry might be produced through turning. From a general point of view, processes can be

classified as: additive (they add material as in metal printing), forming (they shape the material as in forging) or subtractive (they eliminate material as in machining).

2. Checks whether the material characteristics are compatible with the listed processes.
3. Analyses quality requirements (geometric tolerances and surface finish) and determines whether the selected processes can meet the specified requirements or if additional finishing processes, such as grinding, are needed.
4. Sequences the processes and schedules the operations. This allows for estimating production times and costs.

Despite being the most common operational method, this procedure has three disadvantages, which are exacerbated when dealing with short or single-part production runs.

2.1.2.2 Resource matching and allocation service

Resource and subservice matching services are responsible for identifying and assigning the most appropriate resources to each (decomposed) subservice. It is critical for aligning demand-side requirements with the capabilities offered by diverse providers across the federated Tec4MaaSEs ecosystem. It receives an extensive breakdown of a requested service, which specifies its separate subservices and their individual requirements. Simultaneously, the services maintain a comprehensive understanding of the capabilities offered by diverse providers across the ecosystem.

Utilizing these two inputs, the decomposition of requested services and the capabilities of the registered providers, the services' principal functionality is to discover and associate with suitable providers for a requested service. It executes a rigorous matching process, ensuring that the expectations of each decomposed subservice are accurately matched by the available capabilities of the different providers. This important alignment assures that user expectations are properly resolved with the resources and capabilities available within the ecosystem.

The result of this component is a collection of identified provider capabilities which correspond to the subservices. This information then acts as a direct input to subsequent optimization services, which further refine these matches to produce the most optimal service compositions.

For VN1, the decomposed service request is provided by the user when creating a new service request and stored in the system, subsequently prompting the initiation of the optimization process. This workflow is still in progress, and it will be finalized during the next steps.

For VN2, the process begins with a service request submitted by a user through a user interface. Upon receiving this request, the system processes associated design files for detailed decomposition. This decomposition yields comprehensive details including product properties, product features, and a sequence of manufacturing processes, encapsulated as an event. This event is subsequently consumed and stored by the system. Following this, the system begins to prepare the necessary data for optimization, and after identifying suitable resources, it triggers the optimization process.

Furthermore, in a manual approach that provides the user more freedom, users are empowered to explore the ecosystem by searching through a comprehensive list of all registered providers. They can then refine this list by applying filters based on criteria like geographical location and the specific manufacturing services offered, eventually selecting a preferred provider from the filtered results to commence negotiation. The functionality is important for VN3 and its utilization of the platform. However, it is available to all consumers of the platform.

2.1.2.3 Task batching service

Task batching serves as a dimensionality reduction step within the optimization model used in the composition phase. By aggregating similar service requests, it enhances scalability and computational efficiency, especially in high-volume, time-sensitive scenarios Batching facilitates faster and more explainable decision-making by reducing dimensionality and grouping requests that share key characteristics. Depending on the available data, batching can be based on various criteria, including time, service type, location, and quality-of-service (QoS) preferences.

The relevance and utility of batching differ significantly across the three value networks (VNs), based on their planning horizons and operational characteristics.

Batching is applicable in VN1, which handles high volumes of similar service requests related to the daily coordination of EB production and distribution across multiple Arçelik factories and partners. The network operates in a low-variety environment, making it well-suited for batching along all four dimensions (temporal, service similarity, geographic, and QoS-based). Applying these techniques reduces the complexity of optimization and supports timely, large-scale scheduling decisions.

In VN2, which supports customized mold production using additive and subtractive manufacturing, batching is selectively applicable. The environment is characterized by low volume and high variety, limiting the effectiveness of broad batching strategies. However, when molds share similar geometries, materials, or processing constraints, temporal or QoS-based batching may be beneficial to improve coordination and reduce idle capacity. The decision to batch here is opportunistic and context-sensitive.

VN3 focuses on long-term procurement and negotiation processes for infrastructure projects. Here, service requests are infrequent, highly customized, and strategic in nature. As such, batching has no meaningful applicability. The optimization challenges in VN3 are not driven by volume or repetition, but rather by high-level trade-off evaluations across suppliers. No dedicated batching analytics are developed for this network beyond the platform's generic horizontal services.

2.1.3 Analytics for Composition Post-Processing

The post-processing phase aims to re-rank alternative supply chains by incorporating additional user-defined quality of service criteria (see *Figure 3*). To ensure effective and transparent decision-making, the platform implements a two-stage evaluation mechanism. First, a filtering phase eliminates compositions that do not meet acceptability thresholds. This is followed by a ranking phase based on the TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) method that is used to prioritize the remaining alternatives. These approaches complement the optimization process by incorporating flexible, user-defined preferences as additional constraints that optimization couldn’t take into consideration, thus ensuring traceability and adaptability to the user needs.

The filtering phase is strictly non-compensatory, which means that a composition (alternative supply chain) is excluded from the ranking phase if and only if it fails any of the defined minimum thresholds. In the filtering phase, the AND logic is applied across all criteria. The only compositions that satisfy all three simultaneously are retained and move to the next phase of ranking, ensuring that inefficient compositions are excluded before any preferential ranking occurs.

In the ranking phase, the remaining compositions are ranked using the TOPSIS method. Each criterion is transformed to a reward (Quality of Service, Provider Preference- the higher the better) or penalty (Number of providers-the lower the better) numerical score and normalized. At this point, users can define and assign weights to reflect their priorities. Especially for the Provider Preference criterion, the preference score is calculated as the ratio of providers positively rated (Preference score = +1) relative to the total number of providers in the particular composition. For each composition a final score, based on its distance from the ideal and anti-ideal solution of the filtered set of compositions is calculated, ensuring that the final ranking and subsequent user choice is based on structured, traceable, and subjective preferences.

The proposed two-step approach reflects common practice in the MCDA literature, where:

a. filtering is used to reduce the computational complexity and the number of alternatives to be ranked (Lamrini et al. 2020),
b. TOPSIS is a widely used method in supply chain management (Hwang and Yoon, 1981)
c. the three criteria that are used in both phases are typical in the MCDA literature (Shahanaghi and Yazdian, 2009; Uygun and Dede, 2016; Azadeh et al., 2017).

The two-phase approach ensures traceability with no detriment to adaptability as new criteria could be added (or removed) for both phases, while maintaining interpretability, which is inherent in MCDA methods. Finally, to support further analysis or integration with external systems, the platform includes a data export function. Filtered and ranked supply chain configurations can be exported in JSON format, enabling users to conduct additional evaluations, share results with collaborators, or incorporate the data into broader decision-making workflows.

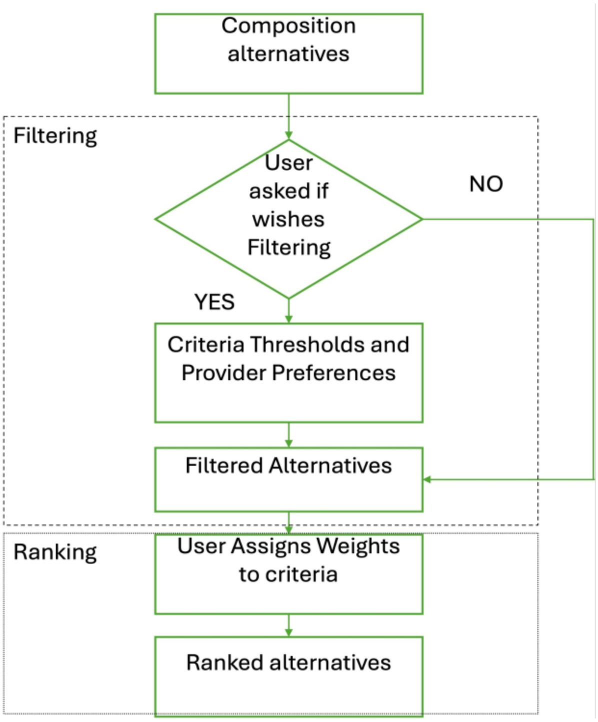

Figure 3: Flow of the interactive filtering and ranking phases

2.1.4 Descriptive analytics and Visualizations

Descriptive analytics and visualizations represent the foundational layer of intelligence within the platform, offering structured insight into manufacturing operations by transforming raw data into actionable knowledge. This component focuses on collecting, aggregating, and interpreting data from a wide range of sources, from user registrations and machine activity to production lines and service interactions, providing a comprehensive view of system performance and status.

Through interactive dashboards, charts, and custom visuals, stakeholders across the value networks can monitor key performance indicators, identify bottlenecks, detect anomalies, and analyze historical trends. These capabilities enhance operational awareness and strategic planning, enabling informed decision-making across all levels of the manufacturing process.

This layer serves as both a reflection of ongoing system behavior and a support mechanism for continuous improvement. Tailored visualizations ensure that each stakeholder, whether consumer, provider, or coordinator, can focus on general information, statistics and metrics most relevant to their role and domain context.

All statistical data is visualized through interactive dashboards built with Power BI, providing a clear, intuitive understanding of system dynamics. The visualization analytics service is divided into three major components, each highlighting a different operational perspective, while maintaining flexibility to accommodate evolving project needs.

1. General statistics reflecting overall platform activity (Section 2.1.4.1)
2. Provider-specific insights for evaluating operational efficiency and resource utilization (Section 2.1.4.2)
3. Post-optimization visualizations and statistics supporting the filtering and final ranking of alternative supply chain configurations (Section 2.1.4.3)

The Visualization analytics layer follows a modular, scalable data pipeline, which includes:

- Data Ingestion: Continuous collection of operational data (mostly in JSON format) from key platform actors (e.g. Marketplace)
- Data Processing: Parsing, transforming, and structuring of raw data for analysis
- Data Storage: Centralized storage (e.g., data warehouse or relational database) optimized for fast, complex querying
- Visualization Layer: Real-time dashboards built in Power BI, connected to the processed data, highlighting trends, KPIs, and system events

This architecture ensures that analytics remain transparent, traceable, and responsive, contributing to the adaptability and efficiency of the platform across multiple value networks.

2.1.4.1 Platform Overview

Various Users (consumers, providers, coordinators), depending on their role, require a centralized, visual interface to track platform activity, including (but not limited to) the status of service requests or the availability and performance of different providers. Without such transparency, stakeholders may struggle to monitor order progress, coordinate effectively, or assess the responsiveness of the network, which can erode trust and limit participation. This component provides a high-level overview of key metrics aggregated across all providers. These include the total number of requests, completion status (e.g., completed, in progress, incomplete), types of services offered, and provider locations—supporting both operational visibility and strategic decision-making.

The dashboard presented in Figure 4 , provides an example of a system snapshot, with features such as:

- Total Requests Overview: Cumulative service requests across the platform
- Request Status: Distribution across completed, in progress, and incomplete requests
- Provider Activity Summary: Request volumes and completion rates per provider
- Service Offerings: Services delivered by each provider (e.g., Milling, Drilling, 3D Printing in VN2)
- Geographic Distribution: Map-based visualization of provider locations

Additional potential KPIs:

- Percentage of requests per service type by status (completed, in-progress, incomplete)
- Request distribution per provider
- Total requests handled per provider

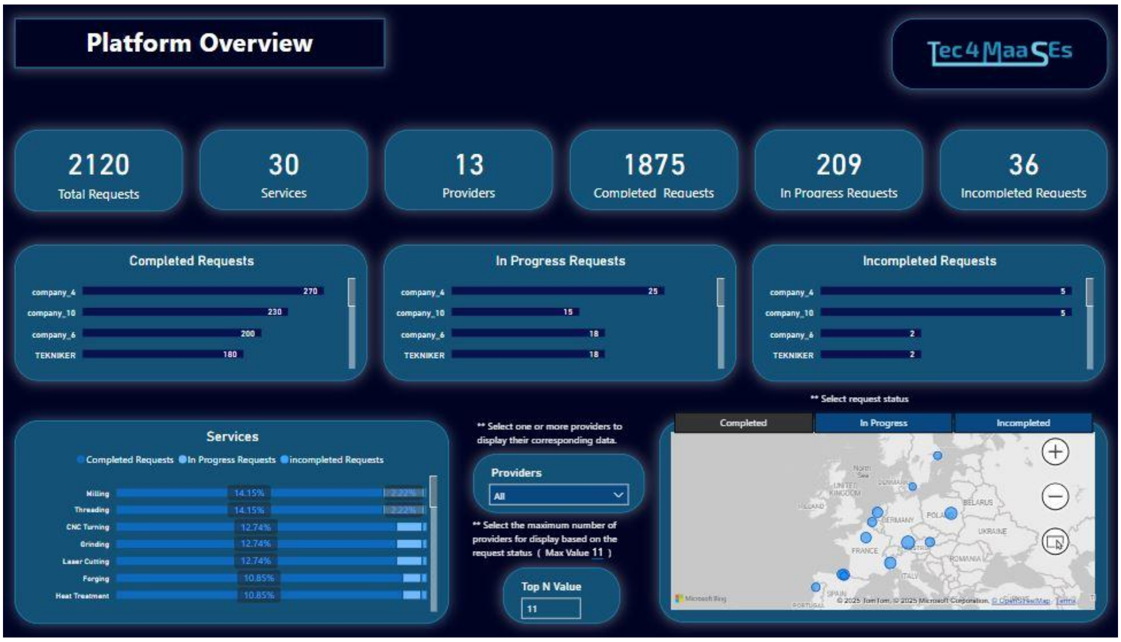

Figure 4: Platform overview dashboard

2.1.4.2 Provider Overview

To maintain high performance in distributed manufacturing, providers need granular visibility into their own operational efficiency—particularly machine usage and request performance. This is especially critical in demanding contexts such as injection mold production and repair, where short lead times and high service quality are essential.

This component addresses these requirements by analyzing provider-specific data, including machine metrics, request outcomes, and service distribution. The dashboard illustrated in Figure 5 allows each provider to track and enhance operational performance through the following features:

- **Request Status Overview**: Total, completed, in-progress, and incomplete requests for each provider
- **Machine Metrics**: Average processing time (in hours) and utilization rate (as a percentage) per machine
- **Service-to-Machine Mapping**: Which services are executed by which machines
- **Operational Efficiency**: Visual indicators highlighting machine operating time, workload balance, and process improvement opportunities

Key Performance Indicators may include:

- User-defined utilization threshold per machines
- Utilization rates by machine
- Percentage breakdown of request outcomes
- Average utilization rate per machine
- Average processing time per machine
- Request distribution by service and status

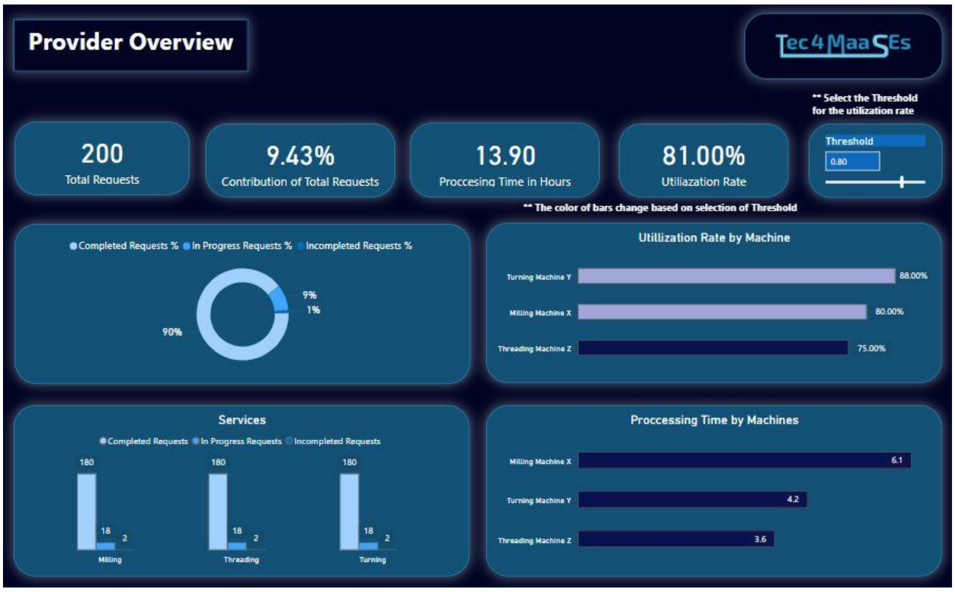

Figure 5: Provider overview dashboard

2.1.4.3 Post-Process Overview

Assessing the outcome of the optimal alternative supply chains is crucial for refining operations and improving future decisions. Post-optimization analysis reveals patterns and performance dimensions not directly addressed during the optimization phase, such as user preferences or subjective quality indicators. This module focuses on evaluating alternative optimal supply chains by aggregating metrics, structural characteristics, and consumer feedback. It supports dynamic analysis of completed request flows, from initiation to delivery, enabling targeted refinements and continuous improvement. Input data includes optimal solutions from the optimization engine, supplemented with user-defined parameters and additional data when available. The dashboard enables interactive features of filtering and ranking of supply chains based on various performance indicators. Further details on the filtering and ranking methods are provided in Section 3.2.5.3.

Suggested Interactive Filtering and Ranking Fields

- **Number of Providers**: Filters supply chains by the number of involved participants to ensure a satisfying level of complexity
- **Quality of Service**: Sets a minimum quality threshold based on aggregated provider scores (e.g., based on delivery delays, defects, number of default deliveries)
- **Provider Preference**: Includes subjective ratings for individual providers (-1: dislike, 0: neutral, 1: preferred)

The system also enables users to assign custom weights to the various ranking criteria, such as the number of providers, service quality, and preference inputs, allowing the ranking algorithm to be adjusted in line with individual strategic priorities. This flexible evaluation method ensures that different users can tailor the decision-support tools to reflect their specific operational goals or quality expectations. The available supply chain configurations are presented in a structured table, which can be sorted and filtered interactively. This facilitates quick comparisons and supports the identification of the most suitable alternatives based on both objective metrics and user-defined preferences. The dashboard illustrated in Figure 6 is a snapshot of the corresponding overview composition interface while Figure 7 shows the drill down functionality of decomposing each value chain into its individual processes and components.

Figure 6. Composition overview dashboard

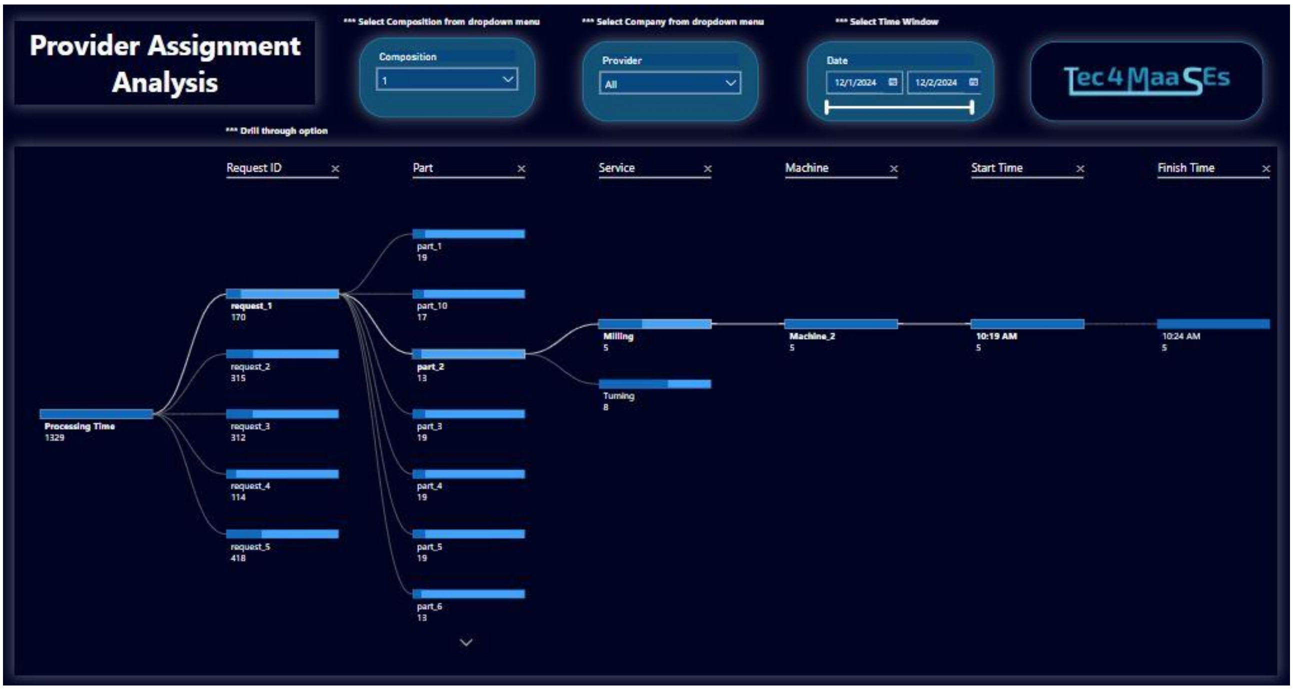

Figure 7. Drill down functionality for requests

Finally, to support further analysis or integration with external systems, the descriptive analytics platform includes a data export function. Filtered and ranked supply chain configurations can be exported in JSON format, enabling users to conduct additional evaluations, share results with collaborators, or incorporate the data into broader decision-making workflows.

3 Preliminary results

3.1 Value Network 1

3.1.1 Background

This section focuses on Value Network 1 (VN1), which comprises four factories involved in the production of white goods and electronic cards. White goods require a wide range of highly customized Electronic Boards (EBs), produced by these factories using more than 500 components, such as PCBs and transistors. These EBs are then supplied to other factories for the assembly of household appliances. Arcelik, a market leader in the European home appliances industry, produces over 30 million appliances annually through a complex value network, including 30 facilities spread across nine countries. In VN1, a significant portion of the required EBs is produced by Arcelik's own factory in Çerkezköy (referred to as Factory 1 or AC). However, due to capacity limitations, outsourcing EB production is sometimes necessary. Karel Electronics (Factory 2 or KE) holds the largest share of Arcelik's outsourced EBs. KE, a leading electronics manufacturer in Turkey, specializes in electronic components and devices, and is renowned for its innovative technology and high-quality products.

On the other hand, the Arctic Romania Washing Machine Factory (Factory 3 or AR) and the Arcelik Bolu Cooking Appliances Factory (Factory 4 or AB) both place orders for electronic boards directly with Arcelik. Since AC cannot meet the total demand for EBs from all factories, including AR and AB, Arcelik has agreements with additional suppliers, such as KE. Under these agreements, each Arcelik factory, acting as a consumer, can independently request the production of the necessary EBs from either AC or KE. The responsibility for producing and delivering EBs to the AR and AB factories is shared between AC and KE.

Production planning at AC is organized once per month through an ERP system that logs inventory and demand data from consumers. The system operates within a flexible time window specified by the providers, allowing for modifications to orders as needed. KE on the other hand, receives quarterly orders from each consumer and allows for their revision at most twice per month. The finalized products are sent to AB daily or weekly via trucks and minivans. In the case of AR, deliveries take place once or twice per week depending on the demand, using large trucks. It is important to note that although both the consumers and one of the providers belong to the same company (Arcelik), in terms of their day-to-day operations they work independently, meaning that each consumer is responsible for their own purchases of EBs and the ordering takes place in a decentralized manner.

The aforementioned processes require integrated synchronization mechanisms, as critical steps are currently performed in a sub-optimal, ad hoc manner, without proper monitoring. This leaves room for potential errors, disruptions across the entire value chain, and inefficient timeframes. On that account, the VN1 goal is to improve coordination and communication with providers in the context of new EBs ordering, production and distribution, by offering their production capacity as a resource to be exploited by the EB consuming plants, in a manner that increases both throughput and resilience.

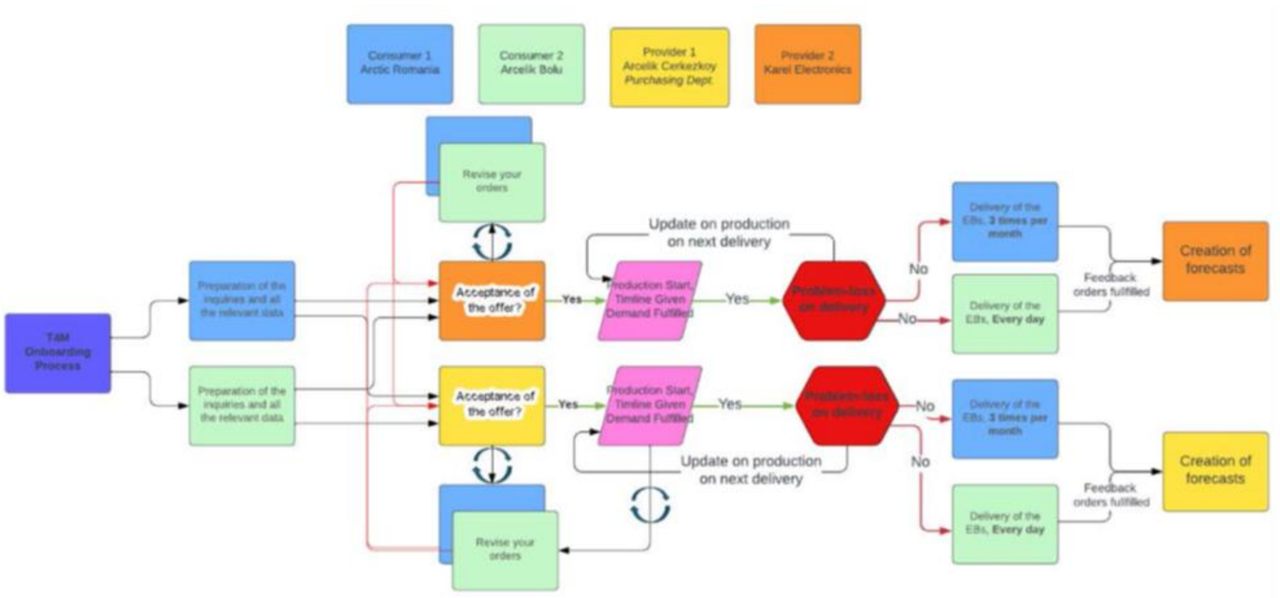

Figure 8: VN1 Phases of the new interactions between different facilities

We assess the data analysis requirements necessary to support the manufacturing and distribution of Electronic Boards (EBs) to meet customer demand in the white goods industry. In this context, we will focus on user stories that require analytics effort specific to the scope of MaaS, identifying the most relevant ones. Additionally, we propose ideas for useful aspects of analytics supported by related literature. Following the discussion in the "Goals" section, we analyze the various stages of VN1 (Figure 8). We conclude that the analytical needs related to VN1 focus on the following key areas (relevant user stories):

User Story 1.4: "As the Planning Department of Consumer (AR, AB) I want a step-by-step wizard which prompts for input data describing product (EBs) and process requirements (e.g. delivery time) because I want to place a request for a manufacturing service". (see Figure 9)

Figure 9. VN1 User case diagram for US1.4

User Story 1.5: "T4M wants to extract the manufacturing service requirements and then match eligible supply chain con-figurations because a ranked list of the supply chain configurations should be returned to the PD". (see Figure 10)

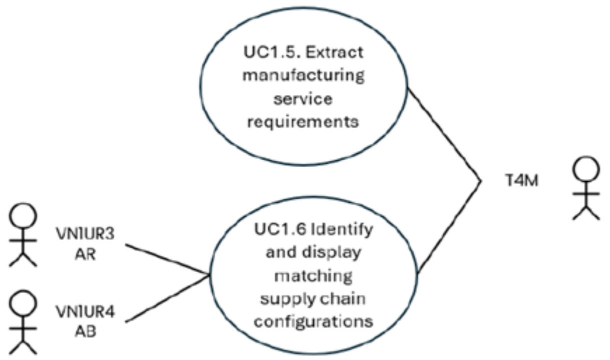

Figure 10. VN1 User case diagram for US1.5

User Story 1.6: "As the procurement representative I want a scoreboard of the proposed supply chain configurations along with a user interface that includes a selection feature, because I want to request service quotations from certain providers". (see Figure 11)

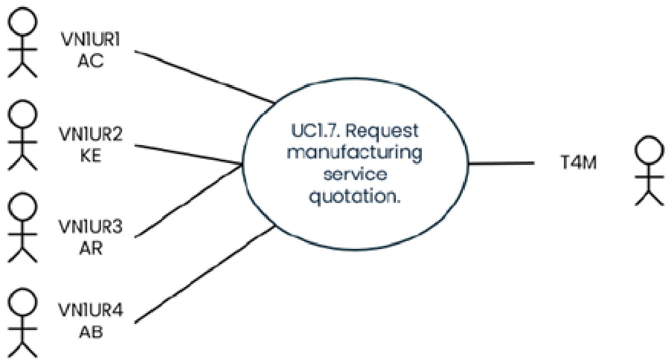

Figure 11. VN1 User case diagram for US1.6

User Story 1.7: "As the Provider Planning Department (AC)/EMS Group (KE) I want a step-by-step wizard to automatically assess my capability to produce EBs by the provided material inventory, production planning schedules and forecasted ETAs to enable me to review requests and release manufacturing services quotations to Customers for EBs I can produce and deliver". (see Figure 12)

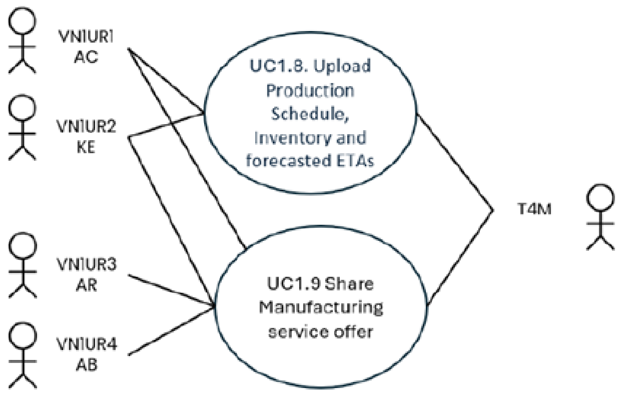

Figure 12. VN1 User case diagram for US1.7

In what follows we provide the list of KPIs as outlined in the DoA.

- **VN1-KPI-1. Production Capacity Utilization Rate**: This KPI measures the efficiency of production by comparing the actual production capacity used against the system's design capacity. The design capacity, recorded in the SAP Production Planning module, serves as the baseline, while actual monthly production is tracked to calculate the ratio. This metric helps assess how effectively the manufacturing system is operating relative to its maximum potential. The objective is to boost the production capacity of utilization rate from 88% to 92% for Arcelik, and from 77% to 90% for Karel.
- **VN1-KPI-2. Order fulfillment rate (OTIF)**: This KPI tracks the percentage of customer orders delivered on time and in full, indicating the company's efficiency in meeting delivery deadlines and providing all requested items. It is calculated by dividing the number of orders delivered on time and in full by the total number of orders placed within a given period. Key data includes order volumes, fulfillment times, inventory levels, and shipping information. Currently, orders are managed via email, while inventory and shipping data come from the SAP WMS system. The baseline performance for Arcelik is 82.91%, with a target of increasing to 95%. For Karel, the baseline is 36%, with a goal of reaching 90%. Achieving these targets will require optimizing inventory management and implementing an order management system.
- **VN1-KPI-3. Material Stocks**: This KPI measures the inventory coverage ratio, expressed in months, to meet current production needs. It is calculated by dividing the current stock quantity by the average monthly consumption. Data for this KPI is available through the SAP ERP system. The baseline coverage is 2.7 months, with the goal of reducing material stock to 1 month for both Arcelik and Karel through improved demand forecasting.
- **VN1-KPI-4. Finished product stock keeping**: This KPI measures the average time finished products remain in inventory before being shipped. The current baseline for Karel is 3 days, with the goal of reducing finished product stock to 1 day through improved demand forecasting. For Arcelik Bolu and Arctic, the baseline is 15 days, with a target of reducing stock levels to 7 days.

Demand Forecasting in Electronic Board Manufacturing:

Traditional forecasting in electronic board manufacturing relies heavily on statistical time series methods due to their proven effectiveness in capturing patterns in high-volume production environments. Hyndman and Athanasopoulos (2018) demonstrate that exponential smoothing and ARIMA models remain highly competitive for manufacturing demand forecasting, particularly when demand patterns are stable with clear trends and seasonality. For electronic component production, Syntetos et al. (2016) found that simple statistical methods often outperform complex machine learning approaches when dealing with regular demand patterns and limited external factors.

The effectiveness of these methods in manufacturing contexts is well-documented. Box et al. (2015) show that ARIMA models excel at capturing both short-term dynamics and long-term trends in production data. Similarly, Gardner and McKenzie (2011) validate that exponential smoothing variants, including Winters' method for seasonal data, provide robust forecasts for production planning horizons of 1-12 months. The rolling-window validation approach recommended by Tashman (2000) has become the standard for time series model selection in industrial applications, ensuring models perform consistently across different forecast horizons.

Capacity Calculation and Resource Allocation:

The literature on capacity management in manufacturing distinguishes between static and dynamic approaches. Traditional capacity planning relies on theoretical calculations based on rated machine speeds and planned operating hours (Hopp and Spearman, 2011). However, these theoretical calculations often overestimate actual capacity by 20-30% due to unaccounted variations in product mix, setup times, and operational disruptions (Vollmann et al., 2005).

Modern capacity management systems have evolved toward real-time, product-specific calculations. Silver et al. (2016) emphasize the importance of maintaining both aggregate capacity measures (for rough-cut planning) and detailed product-based calculations (for finite scheduling). The dual approach—using simplified calculations for new products and learning from actual production data—is supported by Nahmias and Olsen (2015), who demonstrate that capacity estimates improve significantly after initial production runs provide empirical cycle times.

The concept of "available to promise" (ATP) capacity, which considers the portion of capacity actually allocated to specific customers or markets, has become crucial in multi-product environments (Stadtler and Kilger, 2008). This aligns with the T4M approach of tracking allocation percentages as the primary dynamic variable, recognizing that technical capacity remains relatively stable while business allocation decisions drive actual availability.

3.1.2 Methodology

3.1.2.1 Forecasting Methodology

The forecasting service implements a sophisticated BOM-based demand derivation approach that provides more accurate predictions than direct EB forecasting. Instead of forecasting EB demand directly, the system forecasts appliance production (washing machines for AR, ovens for AB) and then derives EB requirements through Bill of Materials (BOM) analysis. This methodology particularly benefits providers (AC and KE) who can align their material procurement and production planning with actual end-product demand patterns. Additionally, consumers (AR and AB) benefit by gaining visibility into their upcoming EB requirements, allowing better procurement timing.

3.1.2.1.1 Forecasting Models

The forecasting framework includes a comprehensive portfolio of statistical models, each suited to different demand characteristics:

ARIMA Models (AutoRegressive Integrated Moving Average) are tested in multiple configurations: (2,1,1), (3,1,1), (3,1,2), (3,2,1), and (6,1,1). These models excel at capturing complex autocorrelation patterns in appliance production data, with the parameters representing: p (number of autoregressive terms capturing how past values influence current values), d (degree of differencing to achieve stationarity), and q (number of moving average terms accounting for past forecast errors).

Exponential Smoothing Family includes three variants. Simple Exponential Smoothing (SES) with smoothing parameters α = 0.2, 0.4, and 0.6 is applied when appliance demand shows no clear trend or seasonality. Holt's Linear Trend method extends SES by adding a trend component, tested with level smoothing (α) values of 0.4, 0.6, and 0.8, and trend smoothing (β) values of 0.2 and 0.4. Winters' Exponential Smoothing (also called Holt-Winters) adds seasonality handling, tested with seasonal periods of 4 and 6 months to capture quarterly and semi-annual patterns in appliance production.

Moving Average Models are included as baseline comparators, using windows of 3, 6, and 12 months. While simpler than other methods, they provide stable forecasts for mature products with consistent demand and serve as a robustness check against overfitting by more complex models.

3.1.2.1.2 Model Validation and Selection Process

The validation employs a rolling window approach with 12-month test periods, evaluating each model's performance at forecast horizons from 1 to 12 months. For each month in the test period, the system trains on all historical data up to that point, generates forecasts for the next 1-12 months, and compares predictions against actual values. Models are ranked by Mean Absolute Percentage Error (MAPE) and Mean Absolute Error (MAE), with the lowest average MAPE across all horizons selected for operational use.

3.1.2.1.3 Implementation Architecture

The service is built using Python's Django REST framework with a PostgreSQL database backend. Key components include REST API endpoints for receiving forecast requests and delivering results, stats models library for ARIMA and exponential smoothing implementations, and PostgreSQL storage for forecast outputs. The BOM explosion process transforms appliance-level forecasts into EB-specific demand projections by multiplying forecasted appliance quantities by the EB requirements specified in each product's BOM.

3.1.2.2 Capacity Calculation Methodology

The capacity calculation system employs a dual-method approach that addresses a fundamental challenge in manufacturing: how to quote capacity for products that have never been produced before, while also leveraging historical data when available. This methodology recognizes that capacity estimates must evolve from rough approximations to precise calculations as production experience accumulates. The dual approach solves the "cold start" problem in capacity planning. When a customer requests a new EB type, providers need to immediately assess whether they can fulfill the order and provide delivery promises. Without any production history, the component-based method provides a reasonable estimate based on the facility's general capabilities. However, once production begins, actual cycle times often differ significantly from theoretical estimates due to product-specific factors like complexity, setup requirements, and learning curves. The time-based method captures these realities, providing the accuracy needed for reliable planning and scheduling.

3.1.2.2.1 Component-Based Capacity Calculation

This method treats all components as equivalent units, using the facility's nominal throughput rate. It answers the question: "Given our facility's general speed and availability, how many components can we produce?"

The calculation follows a hierarchical parameter structure:

Dynamic Parameter - Allocation Percentage (updated weekly): This represents the percentage of total facility capacity allocated to T4M EB production, typically ranging from 40-80%. This is the most frequently changed parameter as it reflects weekly business decisions about capacity commitment to different customers and projects.

Semi-Static Parameters (updated as needed):

- Days/Week: The number of days the production facility operates per week
- Shifts/Day: Number of production shifts operated daily
- Hours/Shift: Standard working hours per shift
- Downtime Percentage: Accounts for all non-productive time including equipment maintenance, shift breaks (meals, breaks), changeovers, and minor stoppages
- Component Per Hour: The facility's nominal production rate, representing how many components the facility can produce per hour under normal conditions. This changes with facility expansions, equipment upgrades, or process improvements

3.1.2.2.2 Time-Based Capacity Calculation

After the first production run, actual cycle times become available, enabling precise capacity calculations. This method answers: "Given how long it actually takes to produce this specific EB, how many can we make in the available time?"

The time-based method captures:

- Product complexity: Simple EBs may take 1.78 minutes while complex ones require 9.03 minutes
- Setup and changeover times: Actual time includes preparation specific to each EB type
- Real constraints: Actual bottlenecks that theoretical calculations miss

3.1.3 Data requirements

The analytics services for VN1 require two distinct data streams: forecasting data that captures the demand patterns through appliance production and BOM structures, and capacity data that reflects the dynamic nature of manufacturing operations. The forecasting service employs a BOM-based approach, requiring historical appliance production data combined with detailed bill of materials to derive EB demand. The capacity calculation service requires both static facility parameters and dynamic allocation decisions, with different data needs for the component-based and time-based methods. All data is currently provided through Excel exports from existing ERP and engineering systems, ensuring security compliance while maintaining data quality. The following tables detail the specific data requirements for each analytics component.

Table 1: BOM Data Requirements

Product	EB SAP code	1921403000
Product Description	EB name	MOTOR KART GR ATLAS32mm VEKTOR 900W
Main Code	Appliance model code	7001440098
Main Code Description	Appliance model name	BEKO WM215
SD	Sub-division/Plant	PO
Quantity	Unit of measure	ADT (Adet/Piece)

Table 2: Appliance Production History Requirements

Product	Appliance model code	7001440098
Date	Production period	202303 (YYYYMM)
ProductionQty	Units produced	5

Table 3: Capacity Calculation Input Parameters

Allocation %	Percentage of capacity for T4M	Weekly
Days/Week	Operating days	As Needed
Shifts/Day	Daily shifts	As Needed
Hours/Shift	Hours per shift	As Needed
Downtime %	Non-productive time	As Needed
Component/Hour	Nominal production rate	When capacity changes

Table 4: Production Time Data (Post-Production)

CODE	EB SAP code	167260304	MES/ERP
DESCRIPTION	EB name	REMIND TOUCH IND 4 GOZ	MES/ERP

Total Time	**Minutes per unit**	**2.96**	**Measured**
EB per hour	**Calculated rate**	**20.2**	**Calculated**

3.1.4 Results

3.1.4.1 Forecasting Results

The BOM-based forecasting methodology has been validated using Turkish automotive industry production data as a benchmark, achieving 70% accuracy (30% MAPE) over a 12-month forecast horizon. This benchmark provides confidence in the approach, as automotive production shares similar characteristics with appliance manufacturing: multiple components per end product, stable but seasonal demand patterns, and BOM-driven component requirements.

3.1.4.2 Implementation Timeline and Expectations

For the initial T4M iteration, the forecasting service will be tested using the actual VN1 data specified in the data requirements—historical production data from Arçelik's washing machines and ovens combined with their detailed BOM structures. Based on the benchmark results and the more stable nature of appliance demand compared to automotive, we expect to achieve at least 70% accuracy (30% MAPE) in this initial implementation.

Further improvements are planned for the final iteration through:

- Advanced model integration: Incorporating more complex models such as Deep Neural Networks (DNNs) for capturing complex non-linear patterns in demand, LASSO regression for automatic feature selection and handling high-dimensional data, and ensemble methods combining statistical and machine learning approaches

- Hyperparameter optimization: Systematic tuning using grid search and Bayesian optimization to find optimal parameters for each model type, including seasonal periods, smoothing coefficients, and ARIMA orders

- External factor integration: Including macroeconomic indicators (GDP growth, consumer confidence), seasonal events (holidays, promotional periods), and market trends that influence appliance demand

We are confident that these refinements will further increase the accuracy beyond the initial implementation. This improvement will directly support the achievement of VN1-KPI-3, enabling providers to reduce material stocks from the current 2.7 months to the targeted 1 month through more accurate demand visibility.

3.1.4.3 Capacity Calculation Results

The dual-method approach provides essential functionality for the T4M Search & Match component, enabling accurate matching between EB service requests and providers with available capacity. This capability ensures that customer orders are matched only with providers capable of fulfilling requirements within specified timeframes. The following analysis presents actual calculations using pilot data and illustrates the system's effectiveness in supporting the Search & Match component.

3.1.4.4 Component-Based Calculation

When a customer submits a service request for a new EB type, the Search & Match component must identify providers with sufficient capacity. Consider a request for 50,000 units of a new EB variant:

- Provider capacity assessment parameters:
- Days/Week: 6 days
- Shifts/Day: 2 shifts
- Hours/Shift: 8 hours
- Downtime: 1 hour per shift
- Component/Hour: 150,000
- Current Allocation: 60%

Calculation process:

- Gross available hours = 6 × 2 × 8 = 96 hours/week
- Total downtime = 2 shifts/day × 6 days × 1 hour = 12 hours/week
- Net productive hours = 96 - 12 = 84 hours/week
- T4M allocated hours = 84 × 0.60 = 50.4 hours/week
- Weekly capacity = 50.4 × 150,000 = 7,560,000 components

The Search & Match component knows that this provider possesses adequate capacity (7.56M > 50K) and includes them in the qualified supplier list.

3.1.4.5 Time-Based Calculation

For established EB types with production history, capacity calculation utilizes actual production rates to ensure precise matching:

Example scenario: Customer request for 1,000 units of PC ARTOUCH2 AK WITH TIMER within 2 weeks

Provider capacity assessment:

- Actual production rate: 9.3 EBs/hour
- Available hours: 50.4 hours/week
- Two-week capacity: 50.4 × 9.3 × 2 = 937 units

Capacity calculation results with insufficient capacity for complete order fulfillment within the requested timeframe.

3.2 Value Network 2

3.2.1 Background

Value Network 2 (VN2) uses Manufacturing-as-a-Service (MaaS) principles to enhance additive manufacturing and machining processes in the injection moulding industry. This value network includes three main stakeholders: ERREKA Plastics, Moldes URA, and Tekniker, bringing their own expertise and capabilities to the network.

ERREKA Plastics specializes in the design and manufacturing of plastic injection components catering to a wide range of industries. Acting as a consumer within the network, it requires innovative moulds tailored to precise design and production specifications. As a provider, it must efficiently manage supply networks and capabilities to deliver manufacturing services, using its available excess capacity. Moldes URA, a small-to-medium enterprise (SME), focuses on the design, feasibility assessment, and production of moulds for different sectors. Their expertise in assessing mould feasibility and preparing manufacturing-ready designs plays a key role in VN2. Tekniker is a leading research center with a state-of-the-art manufacturing shop floor, providing Additive Manufacturing Services (AMS) and Machining Services (MachS) that allow VN2 to explore innovative production methods.

The main goal of VN2 is to simplify and reconfigure value chains in the injection moulding sector by using advanced manufacturing technologies, like additive manufacturing (AM). AM offers many advantages, such as reduced lead times for mould production, enabling faster time-to-market, cost efficiencies by avoiding large spare part inventories, and greater flexibility to meet custom or on-demand production needs. However, there are also challenges, like the high initial investment cost of AM technologies and the need for post-processing and finishing, which can add extra time and cost to the whole production process.

T4M, through VN2, aims to showcase the development of a flexible MaaS marketplace to address these challenges and promote collaboration between providers and consumers. This marketplace presumably will enable (a) dynamic reconfiguration of supply chains to respond to real-time changes in capacity, demand, and logistics; (b) better resource utilization, ensuring that underused assets are offered as services to others, increasing overall circularity; and (c) data-driven decision-making supported by advanced analytics to efficiently match supply with demand. By using the T4M platform, VN2 aims to employ real-time analytics to optimize provider and service selection, apply predictive models to forecast demand fluctuations and anomalies to adapt production schedules, and promote sustainability and resource efficiency through shared production capacities.

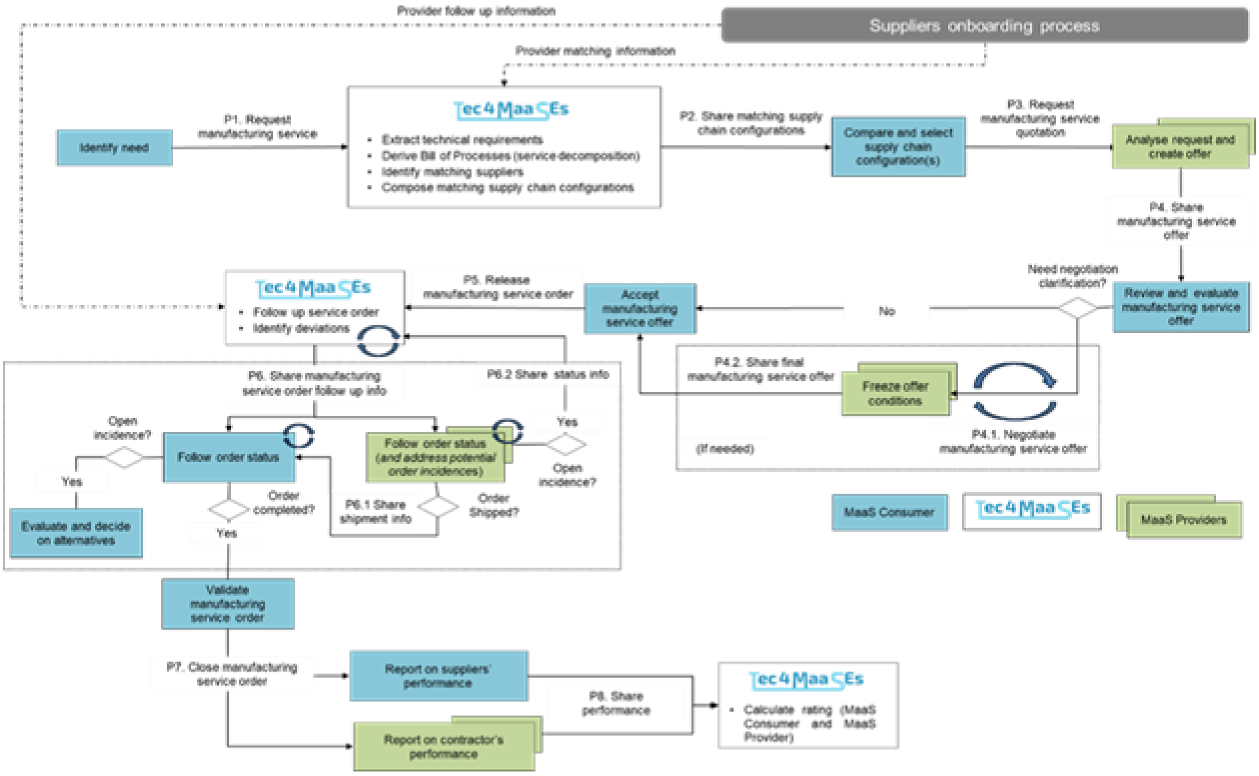

Figure 13: VN2 Phases of the T4M ecosystem. On demand procurement process

Figure 13 outlines the 8 phases and interactions within the VN2 ecosystem, serving as the foundation for the user stories detailed in D2.1. Among these, we identified that the majority of analytics requirements relevant to T4M are primarily associated with user stories US2.P2 to US2.P6. Figure 14 to Figure 18 illustrate the corresponding use case diagrams.

<u>US2.P2</u>: "T4M wants to extract the manufacturing service requirements and then match eligible supply chain configurations because a ranked list of the supply chain configurations should be returned to the procurement representative".

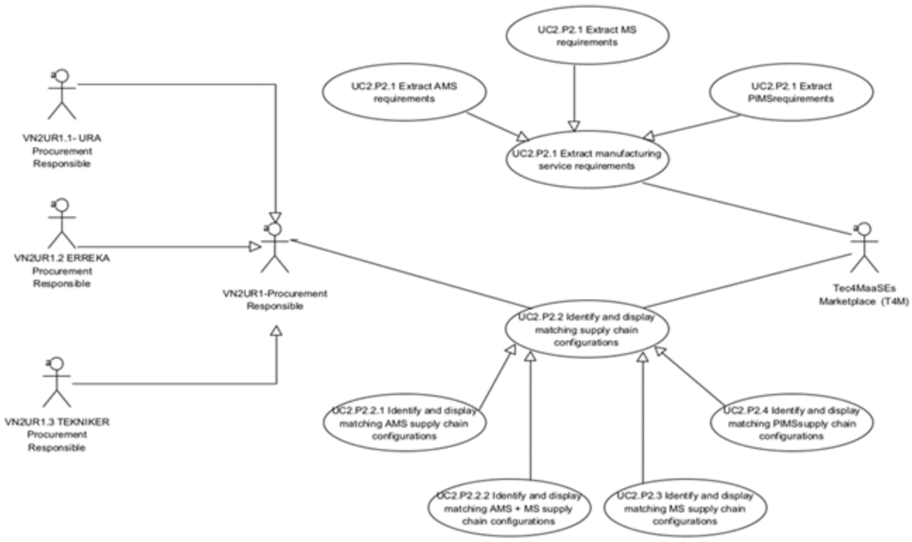

Figure 14: VN2 Use Case diagram for US2.P2

Descriptive and predictive analytics are needed to assist the ranking (optimization) process.

<u>US2.P3</u>: *"As the procurement representative I want a scoreboard of the proposed supply chain configurations along with a user interface that includes a selection feature, because I want to request service quotations from certain providers."*

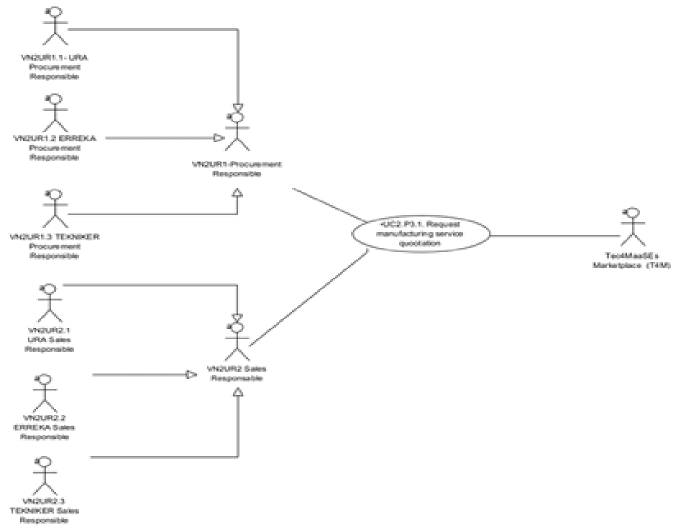

Figure 15: VN2 Use Case diagram for US2.P3

Descriptive exploratory analysis should be in place in order to provide necessary information.

<u>US2.P4</u>: *"As a sales representative, I want a step-by-step offer wizard that allows for the review of the requests for services to release service quotations."*

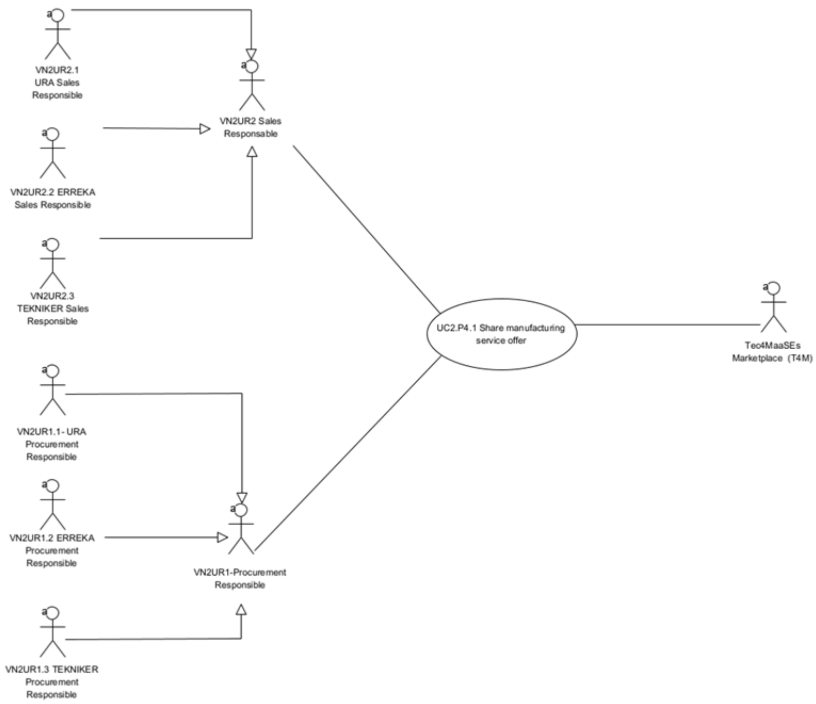

Figure 16: VN2 Use Case diagram for US2.P4

Apart from Descriptive and Diagnostic analytics, also Predictive methods are needed to help estimate appropriate quotes based on past data.

<u>US2.P5</u>: "As a Consumer/Provider, I want a facilitator for the exchange of information to review, negotiate, and release the selected manufacturing service order."

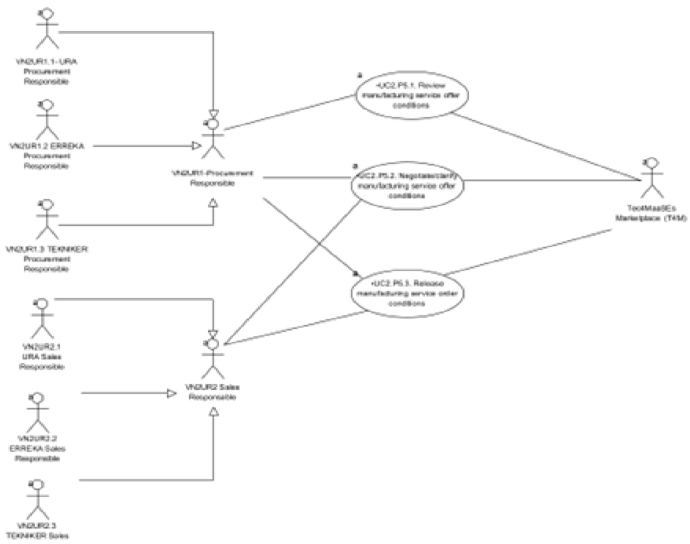

Figure 17: VN2 Use Case diagram for US2.P5

Potential negotiation takes place in this user story calling for methodologies of natural language processing or similarity/deviation detection (see also VN3 for a similar discussion).

<u>US2.P6</u>: "As a Consumer/Provider I want a facilitator for the follow up of the released order because I need to monitor and tackle potential deviations from the original order agreement."

Figure 18: VN2 Use Case diagram for US2.P6

Somehow, the platform should provide some functionalities to detect probable deviations using for example machine learning or NLP.

In what follows we provide the list of the relevant KPIs as outlined in the DoA.

1 **VN2-KPI-1, Machine Capacity Utilization Rate**: Machine Capacity Utilization Rate measures the extent to which a company's machinery is being utilized compared to its full potential during a given period. It is especially crucial in manufacturing environments where machinery plays a significant role in operations. By highlighting periods when machines are underutilized, this KPI helps identify inefficiencies that can be addressed to improve productivity. The KPI is expressed as a percentage, showing the proportion of the machines' available time that is spent in operation.

2 **VN2-KPI-2, Reduction in investment costs (Ownership vs service)**: When focusing on the procurement of manufacturing services, Product Lead Time specifically measures the duration taken from the moment an order is placed with a third-party service provider to the point when the outsourced service or component is received and ready for use or integration into the final product.

3 **VN2-KPI-3, Investment cost in stored part**: This KPI measures the percentage reduction in investment costs when a business transitions from owning manufacturing assets to utilizing a service-based model. It evaluates the cost savings achieved by comparing the total cost of ownership (TCO) of assets with the costs associated to on demand access to manufacturing services provided by third parties. This helps organizations understand the financial benefits of adopting service-based solutions over maintaining and owning their own infrastructure or equipment.

4 **VN2-KPI-4, Product lead time**: This KPI evaluates the effectiveness of on-demand manufacturing in reducing the cost of storing spare parts. On-demand manufacturing aims to produce goods as they are needed, reducing the need for large inventories and, consequently, lowering storage costs.

3.2.2 State-of-art

An analysis of the four KPIs reveals that they are all closely linked to task decomposition, production process selection, supply chain configurations, service requests and offers as well as validation and feedback. They rely not only on real-time descriptive/visualization and predictive analytics but also play a critical role as inputs to optimization processes. Proactive and predictive analytics, in this context, will be used to collect and process relevant data to measure machine utilization (KPI-1), assess potential reductions in investment costs (KPI-2), monitor stored part stock levels (KPI-3), and evaluate finished product lead time in relation to the cost of storing spare parts (KPI-4). Also, other relevant metrics can also be used as proxies to the above 4 main KPIs. Unlike VN1, the primary challenge in this value network lies in estimating those KPIs and relevant metrics within a highly dynamic and/or multifaceted environment. This setting often presents numerous alternative solutions (potential supply chain configurations), making it challenging to filter out and identify the most efficient options, even with the capabilities of a MaaS platform.

Micro operational level; Requests, decomposition and supply chain configurations:

In Value Network 2 (VN2), analytics play a crucial role in supporting KPIs and achieving operational goals during both the onboarding of providers and the dynamic configuration of supply chains. In Phase 1, tools related to information systems and data cleaning and handling are essential for facilitating the integration and validation of provider manufacturing capabilities by processing large datasets and clustering groups of similar capabilities. Machine learning frameworks further enhance this process by extracting and classifying provider capabilities, improving matching accuracy.

In Phase 2, the decomposition of manufacturing requests into subservices calls for the modularization of each request and thus it might require analytics tools to model production workflows and identify the most efficient configurations. Process mining can be helpful for analyzing historical data and tacking bottlenecks, inefficiencies, and dependencies in the current manufacturing processes. Clustering then might be used for grouping together similar tasks based on certain characteristics. Certainly, the breadth of applicability of such techniques is domain specific and depends on what degrees of freedom the consumer requests might have for the allowance of different decompositions.

Predictive analytics techniques can anticipate capacity constraints or potential disruptions, while they can produce suitable data for prescriptive analytics enabling the evaluation and ranking of supply chain configurations. These tools ensure that supply chains adapt dynamically to real-time manufacturing demands, thus optimizing decision-making throughout the value network. In this context, KPI-1 focuses on assessing how effectively machinery is utilized, identifying bottlenecks, and minimizing downtime. Operational-level analytics provide valuable insights into the efficiency of production processes and overall operations (Kutucuoglu et al., 2001; Azwir et al., 2021). Accurately measuring and predicting capacity utilization is critical in VN2, where wide product customization is a key characteristic. This helps forecast the capacity required to meet future demand, ensuring optimal allocation of production resources and early identification of bottlenecks (Rimo & Tin, 2017; Shahidul et al., 2013). Equipment breakdowns and product defects, which are common challenges in manufacturing, have a significant impact on productivity and capacity. Monitoring these events effectively through dashboards and exploratory data analysis facilitates early detection and proactive solutions, fostering overall efficiency.

Additionally, the second phase involves presenting a set of suitable providers ranked based on the identified selection criteria. These providers are structured into supply chain configurations, which may consist of one or multiple providers depending on the outcomes of the matching process. The primary objective is to identify supply chains capable of fulfilling customer requirements and expectations concerning quality, price, delivery time, reliability, and sustainability. In this context, Multi-Criteria Decision-Making (MCDM) is a fundamental tool, offering a structured methodology for analyzing and prioritizing alternatives across a wide range of applications (Bronja et al., 2011).

Service composition plays a critical role in the second phase, as it enables the selection of optimal service combinations from the candidate options for each subtask. Although this is a clear optimization task the data generated by proactive and descriptive analytics is critical for intelligent matching processes utilizing approaches such as constraint programming or metaheuristic methods (Lu & Xu, 2017; Zilci et al., 2015; Tao et al., 2008).

Request of Service quotations and service offers release:

Analytics needs in Phases 3 and 4 are heavily focused on supporting tasks for generating service quotations and releasing manufacturing offers, with the aim of achieving KPIs 2 and 3. In Phase 3, advanced analytics tools capable of processing large datasets are critical for handling manufacturing requests. These tools enable real-time processing of customer demands, production parameters, and historical data. Models, such as regression or random forests, can be used to predict accurate pricing and lead times by analyzing patterns in historical trends and resource availability. Optimization techniques like linear programming and mixed-integer programming in the scope of Task 3.4 benefit from the previous methods since the data produced can then be used in determining cost-efficient configurations while ensuring delivery deadlines and maintaining quality standards.

Although Phase 4 emphasizes on prescriptive analytics methods for evaluating and ranking competing offers data visualization platforms are essential for enhancing transparency by presenting offer comparisons, thus enabling stakeholders (consumers and providers) to make well-informed decisions based on optimization methods.

Furthermore, tracking the final cost of a product and its delivery for selection in matching and forecasting processes requires analytics to identify inefficiencies, helping to minimize operational costs (Patel & Shah, 2014). Due to this reason, we need to define fitness and properness criteria for eligible machines. The transition from conventional mass production to batch production and ultimately to Manufacturing as a Service (MaaS), adds complexity to planning and scheduling and to manage this environment effectively, providers must focus on reducing work-in-progress, eliminating non-value-added operations, and fine-tuning processes. Typical Operations Management techniques such as lean manufacturing and other supply chain approaches are valuable in shortening cycle times, reducing costs, and increasing throughput (although these might be out of scope for T4M). This, in turn, drives higher profit margins and reduces operating expenses (Sahni & Vagrecha, 2022).

Process and Logistics Estimation for Resource Matching:

Since, the real process time in machining is very difficult to calculate, a raw estimation is needed. At least for the initial planning, the process time will depend on the technical properties of the material that should be extracted and also on the characteristics of the machine. Thus, the process time will be dependent on the specific characteristic of the machine (e.g., a machine with high-speed spindle will take less time). In order to provide a more accurate prediction of these parameters, statistical methods need to be applied (e.g., provide a time window, that is a lower and an upper limit on the processing time).

Supply Chain, Logistics, follow up and validation:

In Phases 5 to 8 of Value Network 2, analytics have a key role in supporting KPIs calculation in tasks like following up on manufacturing service orders, validating these orders, and sharing performance data. For Phase 5, it is essential to use real-time monitoring tools that are vital for tracking the progress of manufacturing service orders. These tools should collect data on the fly about production status, delivery times, and any potential delays, disruptions or anomalies, so that stakeholders can easily spot bottlenecks and make the necessary mitigations on the supply chain and logistics process. Predictive analytics, such as time series analysis or neural networks, can also work here, by forecasting disruptions or delays based on past data trends. Since the fourth KPI emphasizes on the importance of predicting and managing inventory levels across suppliers and providers to avoid shortages or excess stocks, this can prove particularly critical at the platform/administrator level, where capacity plays a key role in matching processes for finalizing product compositions based on both capabilities and available resources. The question of assuring transparency among participants (consumers and providers), to manage capacity effectively becomes crucial to prevent disruptions in the supply chain.

In Phase 6, in validating manufacturing service orders, tools for ensuring data integrity and validation are needed to confirm that the service orders meet the agreed specifications. Rule-based systems can also be useful, automating the verification of various order characteristics like quality, quantity, and timelines, making the whole process more reliable.

Phase 7 focuses on sharing performance metrics. Here, data visualization methods (exploratory, diagnostic) are critical for presenting metrics such as on-time delivery rates (related to KPI-2) and production efficiency (related to KPI-1) in a comprehensive format. These tools allow providers and consumers to jointly review performance, identify areas for improvement, and somehow agree on corrective actions if needed. Additionally, blockchain-based systems can enhance trust and transparency by securely recording and sharing performance data across stakeholders.

Finally, in Phase 8, performance sharing and feedback loops can be supported by using sentiment analysis or more generalized natural language processing tools. These tools can help analyse feedback from all stakeholders to find recurring problems and anomalies or explore new opportunities to improve the manufacturing process.

Overall, these analytics tools and methods will help maintain collaboration between providers and consumers, ensuring operational efficiency and achieving KPI improvements across the whole value network. Last but not least, we need to address explainability. For this reason, the Tec4Masses platform will provide analytical scores of the different configurations / compositions.

3.2.3 Methodology

3.2.3.1 Bill of process generator: Requirements extraction and manufacturing process decomposition methodology

The Bill of Processes Generator is designed to support the requirements extraction and manufacturing process decomposition of additive manufacturing and machining process (UC2.P2.1). Its primary function is to analyse 3D CAD files to identify the necessary sequence of additive manufacturing and machining subprocesses, along with the key part properties and requirements that must be considered for selecting suitable manufacturing resources. By leveraging the 3D CAD model, material specifications (e.g., material type, density, explosion risk, etc.) and additional data sources (e.g., quality requirements or user-defined tolerances), the system extracts essential part characteristics—such as dimensions, volume, weight, and geometric features like holes, symmetry, or surface complexity. Based on this analysis, the generator decomposes the requested manufacturing task into a structured Bill of Processes, outlining a sequence of required operations (e.g., additive manufacturing, cutting, or finishing steps). Each operation is associated with relevant capabilities (e.g., milling, turning, drilling, threading, grinding) and process-specific requirements (e.g., minimum number of axes, spindle speed, achievable tolerances), facilitating accurate resource matching and process planning.

The intelligent decomposition of the manufacturing process, covering both additive manufacturing and machining, enables process planning and efficient resource allocation. It also ensures that the most suitable machines and operations are selected based on the part's technical specifications and geometric features. The final decomposition methodology relies on a combination of geometric analysis and domain expert knowledge, to support context-aware manufacturing decision-making.

Additive Manufacturing

The process for extracting requirements and deriving a bill of processes for additive manufacturing begins with the input of two key data sources; that is a 3D CAD file (in formats such as STEP or STL) along with material data. These inputs serve as the foundation for extracting the technical requirements of the part to be manufactured.

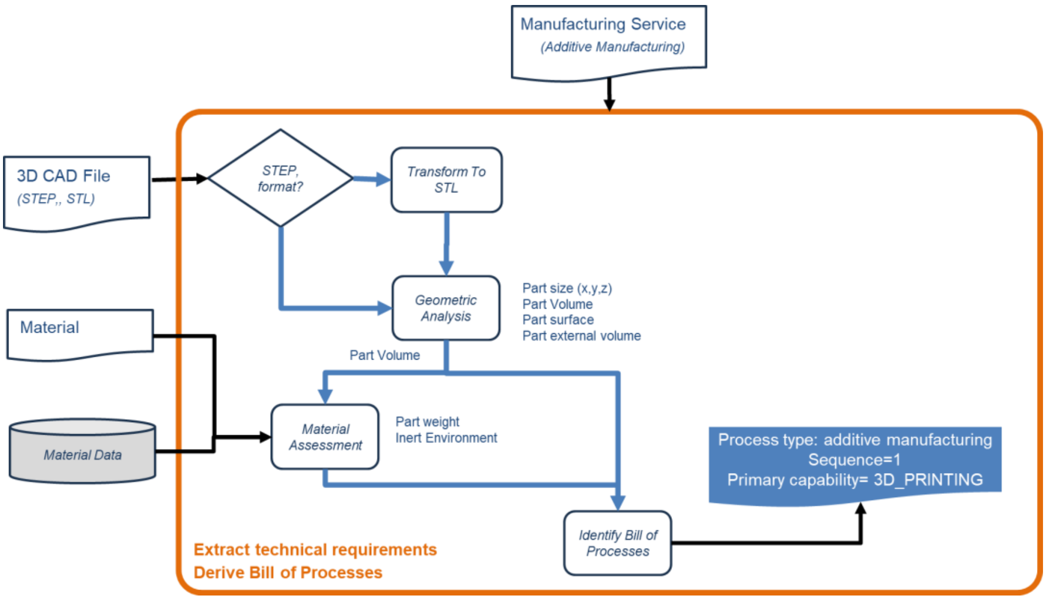

Figure 19: Requirements extraction and decomposition for Additive Manufacturing

The workflow (Figure 19) starts by examining the format of the 3D CAD file. If the file is in STEP format, it is first transformed into STL format, which is more suitable for additive manufacturing applications. This STL file is then subjected to geometric analysis to extract important characteristics such as part volume, which is a key input for subsequent steps. The material assessment process is carried out using the provided material data. This step evaluates important factors such as part weight and the need for an inert environment, which are essential for certain additive manufacturing techniques, particularly those involving metal powders or reactive materials. The results from both geometric analysis and material assessment are then combined to identify the most appropriate Bill of Processes. This includes determining the process type (in this case, additive manufacturing), the operation sequence (starting at sequence = 1), and the primary capability required, which is identified here as 3D printing. If the printed part also requires machining for finishing purposes, the process will then follow the workflow outlined in the following section.

Machining

The requirements extraction and decomposition process for machining starts with the input of a 3D CAD file (in either STEP or STL format) and relevant material data. These inputs are used to extract technical requirements and define the most suitable manufacturing steps for a given part.

For simplification and improved readability, the process is divided into two main stages:

- **Stage 1:** Focuses on the identification of parts, technical constraints, and the derivation of bill of processes by analysing geometric features and assessing material properties to define appropriate process constraints.
- **Stage 2:** Involves extracting technical requirements related to quality constraints and generating bill of processes for machining and finishing operations, such as threading and grinding.

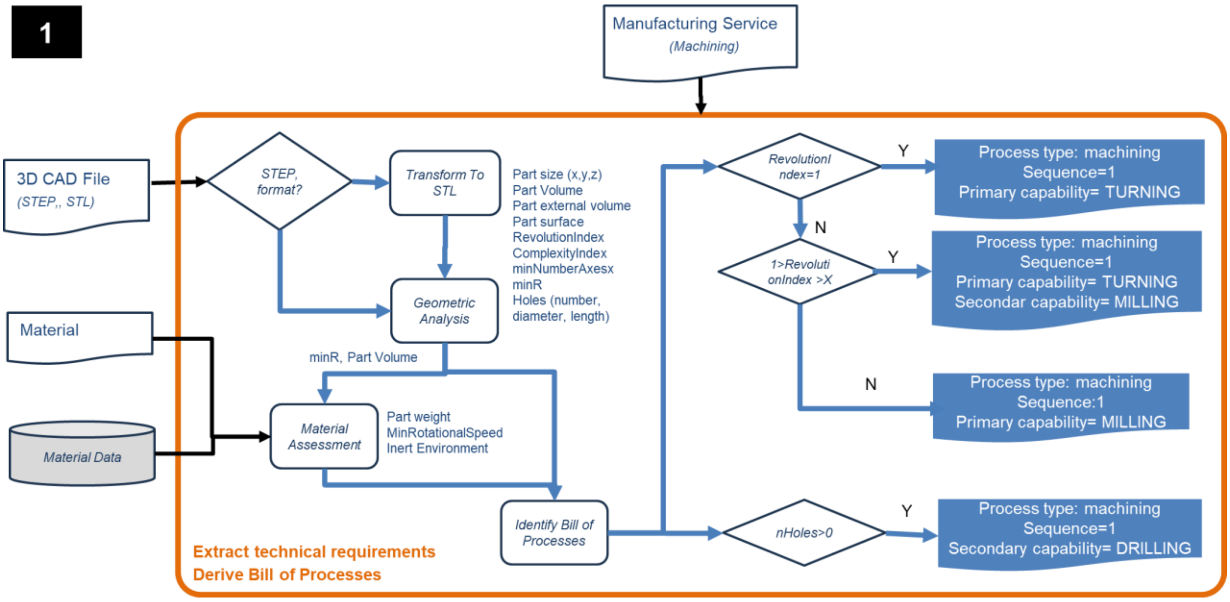

Figure 20: Requirements extraction and decomposition for machining -Stage 1

The workflow begins by checking whether the CAD file is in STEP format. If it is, the file is transformed into STL format to facilitate further geometric analysis. This analysis extracts a range of geometric features, including: part dimensions (x, y, z), volume, visible area and external volume, revolution Index (indicating the degree of symmetry around an axis), complexity Index (indicating the degree of complexity of the part considering elements such as the variety of surface normal and influencing the number of required axes), minimum number of axes required for machining, minimum radius (either for holes, and rounded corners or fillets) and presence and characteristics of holes (number, diameter, length).

Then, a material assessment is performed using the part's volume and minimum radius. This step determines key factors like part weight, minimum required rotational speed and Inert environment requirements (for materials that react with air).

After combining geometric and material data, the system proceeds to identify the appropriate bill of processes. This involves analysing the revolution index and the presence of holes to determine which machining capabilities are required.

- If the Revolution Index = 1, the system identifies Turning as the primary machining capability.

- If the Revolution Index > 1 but below a certain threshold (X), the process involves both Turning and Milling.
- If the part does not meet the turning criteria, Milling becomes the primary capability.
- If holes are present in the geometry, Drilling is added as a secondary capability to the machining sequence.

The following tables present the algorithms developed for the automated geometric analysis of 3D CAD models. These algorithms identify key features of a part to define the required subtractive manufacturing operations. Specifically, they evaluate the degree of geometric revolution to distinguish between turning and milling processes, estimate geometric complexity to determine the need for 3-axis or 5-axis machining, and quantify the presence and distribution of holes to inform potential drilling operations. Additionally, they calculate maximum part dimensions, estimate the blank's weight, and determine the free surface area. Each analysis is based on dedicated analytical models, which are detailed in the corresponding tables.

Table 5: Free surface calculation algorithm

Objective	to compute the total free surface area of a 3D object represented by an STL mesh.
General method	the function calculates the surface area by summing the areas of all individual triangular faces that comprise the STL mesh. Each triangle's area is computed analytically using the cross product of its edge vectors.
Calculated indicators	the algorithm returns the total surface area of the STL mesh, expressed in square millimetres (mm^2). This quantifies the combined area of all mesh faces.
Expected output	a scalar value of the total free surface area [mm^2].
Algorithm description:	1. Data acquisition: the STL file is selected interactively. Once selected, the mesh is read and stored as a list of vertex coordinates and triangular face indices. 2. Surface area calculation: the function iterates through all triangular faces in the mesh: • For each triangle, the three corner vertices are extracted using the face index list. • Two edge vectors are formed by subtracting vertex coordinates. • The area of the triangle is computed as half the magnitude of the cross product of the edge vectors. Triangle area accumulation: all triangle areas are accumulated to obtain the total surface area.

Table 6: Revolution symmetry detection algorithm (turning vs. milling)

Objective	to determine whether the geometry of a CAD model exhibits revolution symmetry around one or more principal axes (X, Y, or Z). This distinction is critical for identifying whether the part is suitable for turning (lathe-based manufacturing) or requires milling operations.
General method	the algorithm performs a statistical analysis of the part's 3D geometry by converting the point cloud into polar coordinates with respect to each principal axis. It quantifies how often radial-height ($\rho \cdot Z$) combinations are repeated, which is indicative of rotational symmetry. This analysis is conducted over three permutations of the axes to ensure all principal directions are evaluated.
Calculated indicators	• Pure revolution index: a binary indicator (1 or 0) that is set to 1 when a significant proportion of ($\rho \cdot Z$) value repetitions exceed a fixed threshold. It indicates clear

	rotational symmetry. • Graded revolution index: A weighted average measuring the extent to which repeated patterns dominate, allowing for near-revolution classification.
Expected output	the algorithm classifies the part as: • Revolution part around one axis (e.g., suitable for turning). • Multi-axis revolution part (e.g., a sphere). • Nearly revolution part, when the geometry closely resembles revolution but with slight deviations. • Non-revolution part, requiring dominant milling operations.
Algorithm description:	1. Data acquisition: the STL file is selected interactively. The part's mesh is parsed to extract its vertices and face connectivity. 2. Axis permutations: the analysis is repeated for three axis permutations (X, Y, Z) as potential revolution axes. 3. Coordinate transformation: for each axis, the vertex positions are converted into polar coordinates: ρ is the radial distance in the orthogonal plane, Z is the axial height and θ the angular coordinate (not used). 4. Discretization and filtering: values of ρ and Z are rounded; if there are two or fewer points at ρ = 0 (cone tips), these are filtered out to avoid noise. Z values are adjusted so all are positive integers. 5. Pattern metric construction: The product ρ·Z is computed for each vertex, then sorted to facilitate pattern detection. 6. Repetition analysis: • Unique values of ρ·Z are identified. • Repetition counts are computed. • Values whose repetition frequency exceeds a defined threshold are flagged. • The pure revolution index is set to 1 if all evaluated values satisfy this threshold condition. Otherwise, it is set to 0. • The graded index is computed as the weighted average repetition frequency of the flagged values. 7. Classification Decision: • Multi-axis revolution: if pure index = 1 in at least two axes. • Single-axis revolution: if pure index = 1 in only one axis. • Near revolution: if no pure index = 1 but graded index exceeds a defined threshold. • Not a revolution: otherwise.

Table 7: Geometric complexity detection algorithm (3-axis vs. 5-axis)

Objective	to assess whether the geometry of a CAD part is better suited for machining with a 3-axis or a 5-axis CNC machine. This decision is based on the diversity and orientation of surface normals, which reflect the geometric complexity and accessibility of the part.
General method	the algorithm analyses the orientation of the surface normals extracted from the STL mesh. A high concentration of similar normal vectors suggests that most surfaces are aligned in common directions, which can be handled by 3-axis machines. A more diverse distribution implies complex surfaces requiring multi-directional tool paths, favouring 5-axis machining.
Calculated indicators	repeatability index. It is a continuous index between 0 and 1 that quantifies the variability of face normal orientations. A lower value suggests higher redundancy (3-axis), while a higher value indicates greater directional variation (5-axis).

Expected output	the algorithm returns one of the following classifications: • 3-axis machining: if the repeatability index is below a defined threshold. • 5-axis machining: if the repeatability index is above or equal to the threshold.
Algorithm description:	1. Data acquisition: the user selects an STL file interactively. The mesh geometry is read and decomposed into vertices and face connectivity. 2. Normal vector computation: using the triangulated mesh, the algorithm calculates the normal vector of each face. 3. Sorting and grouping normals: • All normal vectors are sorted row-wise. • The algorithm identifies sequences of adjacent normals that share at least one common component. • These sequences are grouped using cumulative indexing to capture consecutive similar normals. 4. Repetition analysis: the size of each group of similar normals is counted. The largest group size is used as an indicator of normal vector uniformity. 5. Repeatability index calculation: the repeatability index is computed by analysing how dominant the most common group of similarly oriented surface normals is compared to the total number of surfaces in the mesh. If a large proportion of the normals share a similar orientation, the index will be low, indicating low surface complexity. Conversely, if the orientations are more varied and no single group dominates, the index will be higher, reflecting a more complex surface topology. 6. Classification decision: • If the repeatability index is below a defined threshold, the part is considered suitable for 3-axis machining. • Otherwise, the part is classified as requiring 5-axis machining, due to the geometric complexity implied by its varied surface normals.

Table 8: Minimum concave radius detection algorithm

Objective	to detect circular and arc-like features, determine their concavity or convexity, and calculate the minimum radius among the concave features. This minimum concave radius, combined with the properties of the material to be machined, is essential for estimating the minimum spindle speed (RPM) the machine must be capable of, in order to machine that internal feature.
General method	The function first identifies and clusters nearby vertices around a given point on the STL mesh, forming local neighbourhoods based on spatial proximity. These groups are analysed to determine whether their geometry approximates a planar circular arc. Each group is projected onto a plane and fitted with a circle using a robust least-squares method with outlier rejection. Valid fits are evaluated based on arc angle and point density. The spatial relationship between the fitted circle and the original 3D geometry is used to classify the arc as concave or convex. Finally, the algorithm selects the smallest radius among all arcs classified as concave and meeting quality criteria.
Calculated indicators	minimum concave radius. The smallest radius of a valid concave circular arc is detected.
Expected output	a single numerical value representing the minimum concave radius detected. If no valid arc is found, the result is empty or undefined.
Algorithm description:	1. Data acquisition: a triangular mesh data—consisting of vertices and face connectivity—is extracted from the STL file. 2. Neighbourhood identification: for each candidate point on the mesh, nearby points within a specified distance are grouped together to form local geometric clusters. This helps isolate potential features like edges or arcs.

	3. Planar grouping: each point cluster is examined to determine whether the included points can be split into one or two approximately planar groups. This step ensures that only geometrically coherent regions are analysed further. 4. 2D projection: points in each planar group are projected onto a local 2D coordinate system that lies on the plane. This simplification allows for more robust circle fitting. 5. Arc detection via circle fitting: the projected 2D points are used to fit a circular arc using a least-squares method. The fitting process includes an outlier filtering step to discard inconsistent points and improve accuracy. 6. Fit evaluation: the arc candidate is evaluated based on two criteria: how closely the points adhere to the fitted circle (geometric tolerance), and how much of the circle is spanned by the arc (angular coverage). 7. Concavity check: if the arc meets the quality criteria, the algorithm determines whether it represents a concave or convex geometric feature. This involves comparing the curvature direction relative to the mesh surface. 8. Radius selection: this process is repeated across the mesh to identify all valid concave arcs. Among them, the smallest radius value is retained as the final indicator of minimum concave feature size.

Table 9: Hole detection and classification

Objective	automatically identify and quantify cylindrical holes in a CAD mesh based on detected circular features. This includes determining their dimensional properties (radius and length) and grouping similar holes for summarised output. This is effective for identifying drilling operations during machining planning.
General method	the algorithm first filters and analyses complete circular features extracted from the STL mesh. It calculates the geometric planes associated with each circle and examines all unique pairs to identify potential holes. Pairs are validated based on plane orientation, centre alignment, radius similarity, and concavity characteristics. Candidate pairs are processed starting from the smallest radii, and a pair is only accepted if its centres are not enclosed within any previously detected hole. Finally, holes are grouped by radius and length to provide
Calculated indicators	the algorithm outputs both the total number of holes detected and a grouped summary detailing hole radius, length, and the number of occurrences for each type.

Expected output	• Total number of distinct cylindrical holes detected in the mesh.
	• A list of grouped hole types, where each group contains:
	    • Hole radius [mm]
	    • Hole length [mm]
	    • Number of occurrences
	
	Coordinates of involved centres (for spatial reference)

Algorithm description:	1. Data Acquisition: the algorithm receives a filtered set of full circular features from the surface mesh (calculated through the *minimum concave radius detection algorithm*). Each circle includes a centre point, radius, concavity information, and a set of 3D points that define its contour. 2. Plane computation: for each circle, the algorithm calculates the normal vector of the best-fitting plane. This is done using three representative points from the circle to determine the plane’s orientation and a point of reference (the circle centre). 3. Pairwise matching and validation: every unique pair of circles is analysed for possible pairing:

- The respective planes are checked for parallelism.
- The spatial alignment between centres is examined to confirm axial alignment with the plane normal.
- The distance between planes is compared to the distance between the centres to ensure consistency.
- Radii are compared to verify dimensional equivalence within a defined tolerance.

4. Hole detection: pairs that satisfy all geometric conditions are classified as valid holes. Their length is taken as the distance between the circle centres, and the radius from the shared dimension of both circles. Before confirming the pair, the algorithm checks that the involved centres are not already spatially enclosed within a previously detected hole (to avoid redundant detection in nested or concentric cases). Once a pair is accepted, the constituent circles are excluded from further consideration.
5. Geometric grouping: all validated holes are grouped based on radius and length. If multiple holes fall within the same dimensional range (within the specified tolerance), they are counted as a single group. Each group stores the representative dimensions and spatial distribution of the holes.
6. Summary generation: a structured output is generated that includes:

- Total number of detected holes.
- Groupings of similar holes with their shared dimensions and counts.
- Locations of all contributing centres, useful for spatial mapping or visual verification.

2

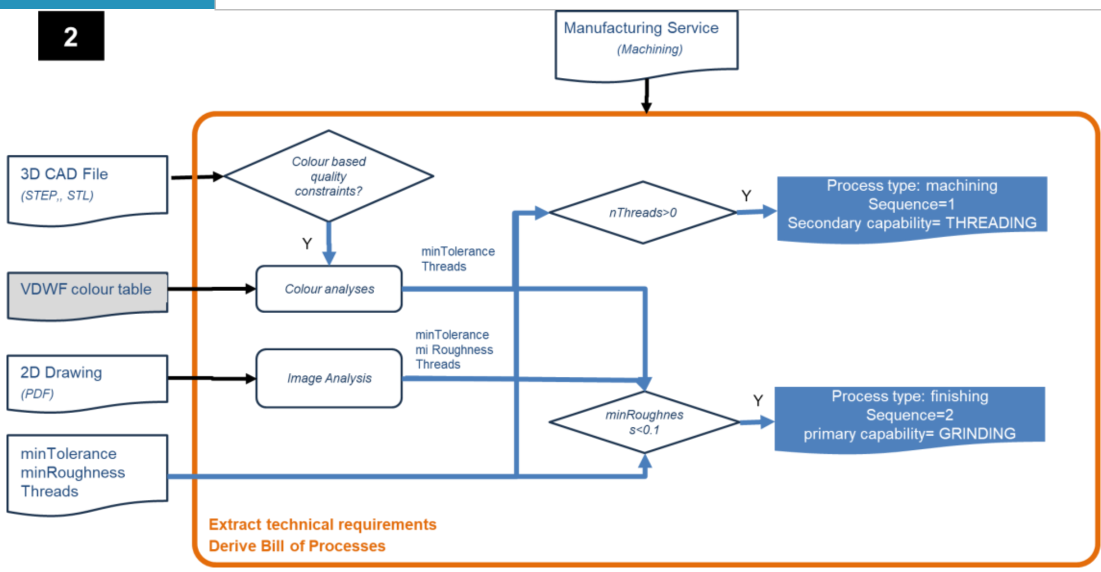

Figure 21: Requirements extraction and decomposition for machining -Stage 2

This workflow focuses on the extraction of technical requirements related to quality constraints and the generation of a Bill of Processes for machining and finishing operations such as threading and grinding.

The process can start from multiple data sources, which vary according to context and the user's internal procedures.

- **Alternative 1:** A 3D CAD file (in STEP format) including colour-based quality constraints based on the VDWF colour table[1], which encodes quality requirements visually in the CAD model in internally stored. The VDWF colour table refers to a standardized colour coding system established by the VDWF (Association of German Tool and Mould Makers). It is primarily used in the tool and mold making industry for annotating and visualizing the status, changes, and features of 3D CAD models during the design and manufacturing process.
- **Alternative 2** A 2D drawing (in PDF format), which includes tolerance specifications, roughness, and threading details.
- **Alternative 3:** Manual data inputs specifying minimum tolerance, surface roughness, and threading information.

In the case of Alternative 1, the workflow first checks whether the CAD model contains color-based quality constraints. If so, a color analysis is performed to extract parameters such as: Minimum tolerance and Thread specifications. For Alternative 2, an image analysis is conducted on the 2D drawing to identify quality-related features, including: Minimum surface roughness, minimum tolerance, and thread specifications. With Alternative 3, these values are entered directly by the user.

Based on the extracted information, if threading features are identified, the system assigns Threading as a secondary capability. Additionally, if the minimum surface roughness requirement is particularly stringent (e.g., *minRoughness* < 0.1), a finishing process is appended, with Grinding designated as the primary capability.

Alternative 1: Colour extraction from STEP files

STEP files allow the possibility of applying colour to the different faces of the 3D model. Although different initiatives have been tried to link colour with tolerances, the VDWF is the one that has attained more diffusion.

[1] vdwf_guidelines_color_table_for_cad_supported_relaying_of_tolerances.docx

Tolerances

R	G	B	Tolerances [mm]	Positions tolerance [mm]
255	217	102	±0,002	
255	175	175	±0,005	
255	128	0	±0,01	
128	128	0	±0,02	
64	255	64	±0,05	
183	183	220	±0,1	
95	0	0	±0,5	
0	0	255	H7	±0,01 when drilling
0	255	255	H5	±0,01 when drilling

Threads

R	G	B	Tolerances [mm] / Specification	Positions tolerance [mm]
255	255	0	Metric standard threads	±0,1
255	175	0	Metric fine threads	±0,1
211	45	96	Imperial threads	±0,1
255	95	0	Non-standard threads	±0,1
0	175	175	Drilling ±0,1	±0,1

Figure 22: Example of the VDWF colour table

This colour table links specific RGB (Red Green Blue) with the numeric values of the tolerances. The algorithm checks the content of the STEP file and try to extract the minimum tolerances and threading values.

Alternative 2: Analysis of 2D drawings

Two different approaches have been taken to extract the information contained in the 2D drawings: (1) OCR based detection and (2) AI based detection which are described below.

OCR based detection

This approach is designed to identify specific elements within a 2D technical drawing using OCR (Optical Character Recognition) technology. Various engineering standards define standardized symbols for features such as surface roughness, threading, and tolerances. The algorithm focuses on detecting these symbols and returning the minimum specified value among them. The key symbols targeted for recognition include:

- ***+-*** : Plus minus sign indicates tolerances. Can be indicated with a single symbol ±, or two individual values + and -, usually one on top of the other.
- ***Ra***: Surface roughness value indicating the smoothness of the surface.
- ***Ø/M***: Used to indicate threading.

The algorithm detects those symbols and returns the minimum value for each of them.

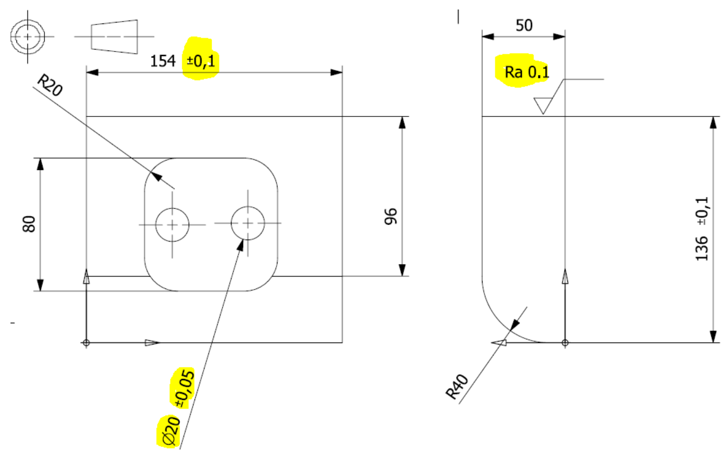

Figure 23: Example of 2D drawing with highlighted points if interest

<u>Use of AI tools</u>

The feasibility of using AI tools to detect surface roughness and dimensional tolerances from 2D technical sketches was explored. Initially, ChatGPT was tested as an automation tool by developing a prototype capable of analysing sketch images to extract tolerances and roughness annotations and outputting the information in a predefined structured format (JSON).

3.2.3.2 Bill of Processes Generator: Processing time estimation methodology

The component provides a rough estimate of the processing time needed to identify the identify the matching supply chain configuration considering available capacity (UC2.P2.1). It is important to note that actual process duration can vary significantly depending on factors such as machine capabilities, the chosen processing strategy, and the specific tooling used. Given this complexity, accurately calculating precise processing times falls outside the scope of Tec4MaaSEs. Instead, the component offers an approximate time estimate sufficient to support initial planning and resource assessment.

This section outlines the methodology and implementation approach used to develop the processing time estimation functionality within the T4M system. The goal was to create a model that provides an initial estimate of manufacturing time.

The estimation logic is designed to be compatible with different manufacturing strategies—subtractive machining (MACHS), additive manufacturing (AMS), or a hybrid of both (AMS + MACHS). The proposed methodology balances simplicity and generality, enabling early-stage decision-making in process planning, while allowing room for integration with more sophisticated simulation or CAM tools for refined predictions.

The processing time is estimated using different formulas depending on the selected manufacturing process.

- **Subtractive machining: (MACHS);** The process time for machining considers the amount of material that has to be eliminated from the raw stock (roughing) plus an additional finishing pass across the surface of the part. This approach estimates the time required for roughing (bulk removal of material) and a single finishing pass.

$$T = (\text{External Volume} - \text{Part Volume}) / MRR_{Roughing} + \text{Part Area} / MRR_{Finishing}$$

- **Additive Manufacturing (AMS)**: Additive manufacturing considers the time that is needed to add the material to form the part.t assumes the part is built directly by material deposition at a uniform rate.

T= Part Volume)/ Deposition Rate

- **Hybrid Manufacturing (AMS + Machining):** If the part requires a finishing machining operation following AMS, given the surface roughness and the dimensional characteristics of the deposited beads, two finishing passes are assumed necessary. The approach reflects the typical need for two finishing passes due to the surface roughness and dimensional variance associated with additive layers.

T= (Part Volume)/ Deposition Rate) + 2 x (Part Area)/ $MRR_{Finishing}$)

Where:

- **Deposition Rate**: Rate of material deposition for additive processes [mm3/s][mm^3/s].
- **Material Removal Rate (MRR) – Roughing**: Volume of material removed per second during bulk material removal [mm3/s][mm^3/s].
- **MRR – Finishing**: Surface-based removal rate for finishing passes [mm2/s][mm^2/s], assuming a single pas **Part Volume**: The actual volume of the final part [mm3][mm^3].
- **External Volume**: The bounding box or cylinder volume enclosing the part [mm3][mm^3], representing the raw stock.
- **Part Area**: Total exposed surface area of the part [mm2][mm^2].s under nominal conditions.

3.2.3.3 Bill of Processes Generator: Implementation

Both manufacturing process decomposition and processing time estimation algorithms are implemented using Python libraries and exposed in a single API. The most relevant libraries are the following:

- ***Flask***: Used to publish a RESTful API to interact
- ***Trimesh/OCC***: Used to load and analyse the CAD files. Also using to transform between STP and STL formats.
- ***Numpy/scipy***: Used for numerical operations in the radius and complexity detection
- ***MinIO***: Used to read from a minIO file storage
- ***Kafka***: Used to publish the results of time-consuming operations
- ***Easyocr***: Used for OCR operations

The components have been tested in a local manner using Postman to call upon the functionalities and minIO’s web browser to upload the CAD files to test.

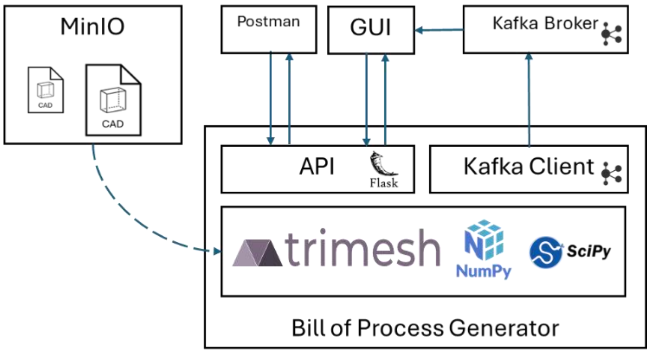

Figure 24: Bill of Processes Generator testing architecture

An initial GUI was also developed using the Panel python library to test the API. It showed a 3D interactive visualization of the CAD model, with the detected radius and holes, as well as panels showing the detected characteristics.

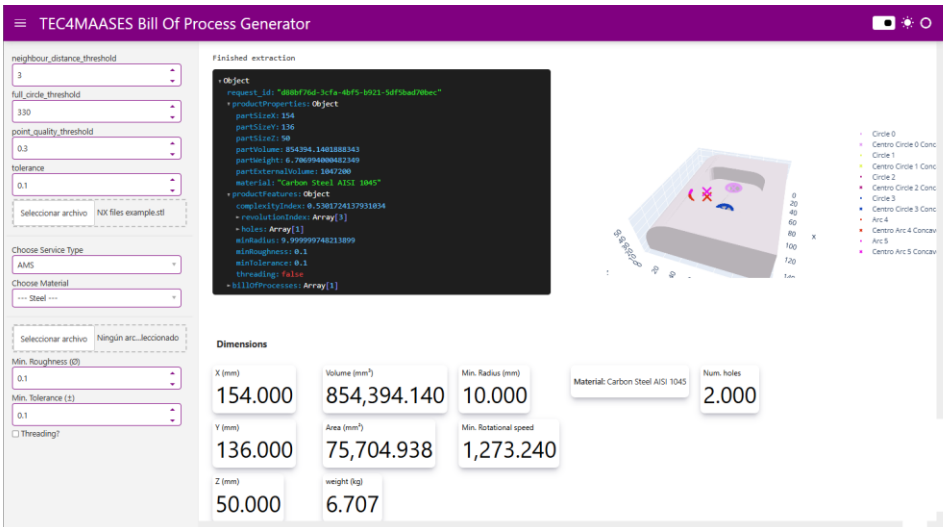

Figure 25: Bill of Process Generator testing GUI

3.2.4 Data requirements

This section summarizes the input data already identified in the previous sections.

3.2.4.1 Bill of Processes Generator: Requirements extraction and manufacturing process decomposition data requirements

To extract requirements and provide process decomposition functionality within the T4M system, a set of core input parameters are required as described in the previous section.

Table 10: Requirements extraction and manufacturing process decomposition input data

Category	Input data	Description	Unit	Source
Workpiece Information	3D CAD File	A 3D CAD (Computer-Aided Design) file is a digital representation of a physical object in STEP or STL format	-	Uploaded by the user

Category	Input data	Description	Unit	Source
	2D Drawing	A pdf file include in the technical drawing of a workpiece.	-	Uploaded by the user
	Minimum tolerance	The smallest allowable variation in a part's dimension from its nominal (intended) size. It defines how much a feature can deviate while still being acceptable.	mm	Defined by the used
	Minimum roughness	The lowest acceptable surface texture or smoothness of a part. It is a measure of the fine irregularities on the surface, often resulting from the manufacturing process.	µm	Defined by the used
Workpiece Material Information	Material type	This refers to the specific kind of material used in the manufacturing process		Selected by the user
	Density	Mass per unit volume of a material,	Kg/m^3	Predefined in a material database
	Minimum Vc	Minimum cutting speed	mm/s	Predefined in a material database

3.2.4.2 Bill of Processes Generator: Processing time estimation data requirements

To implement the processing time estimation functionality within the T4M system, a set of core input parameters are required as described in the previous section.

Table 11: Processing time estimation: Input data

Category	Input data	Description	Unit	Source
Machine Information	Deposition Rate:	Material deposition rate for additive processes. Reflects volume added per second.	mm^3/s	Manufacturing resources specifications
Workpiece Material Information	MRR – Roughing	Material removal rate during bulk roughing operations in machining.	mm^3/s	Predefined in a material database
	MRR – Roughing	Material removal rate during bulk roughing operations in machining.	mm^3/s	
	MRR – Finishing	Material removal rate during surface finishing, modelled per surface area.	mm^2/s	
3D Model Geometry	Part Volume	Volume of the final part derived from the CAD/3D model.	mm^3	Derived automatically from the 3D CAD model using geometric
	External Volume	Bounding volume of the raw stock	mm^3	

Category	Input data	Description	Unit	Source
		from which the part is machined.		processing algorithms.
	Part Area	Total exposed surface area of the finished part.	mm^2	

3.2.5 Results

This section presents the outcomes of initial validation studies conducted during the development phase to assess the performance of the proposed algorithms involved in the Bill Of Processes Generator. Two key areas are addressed: (1) requirements extraction and manufacturing process decomposition, which focuses on the capability of the algorithms to accurately identify geometric and machining features of parts, and (2) machining processing time estimation, which evaluates the precision of the developed time prediction tool against established CAM program benchmarks. For each area, representative data sets were analysed, and results were compared with expert evaluations or theoretical machining times to determine consistency, accuracy, and potential areas for improvement. In addition to the core validation of Bill of Process Generator, this section also presents results from the post-optimization phase, where the various generated compositions were filtered and/or ranked to identify the most effective alternative. This complementary analysis not only illustrates the robustness of the analytics framework but also provides a deeper understanding of its practical implications in real-world MaaS systems. Collectively, these findings provide a comprehensive system's overview, highlighting both its current strengths and limitations for further enhancement.

3.2.5.1 Bill of Processes Generator: Requirements extraction and manufacturing process decomposition results

To verify the correct functioning of the algorithms described in Section 3.2.3.1 for process decomposition, a series of test were performed using a set of representative parts were tested.

Initial results for extraction and manufacturing process decomposition: machining

For each case, the output produced when selected some CAD files was compared against the expected result to assess consistency and accuracy. The following table summarises the parts analysed, the results obtained, and whether they align with the assessment of a domain expert.

Table 12: Initial results for extraction and manufacturing process decomposition: machining

Part (STL view)	Output from algorithms				Accordance with expert
	Revolution symmetry	Geometric complexity	Minimum complex radius	Holes and classifications	
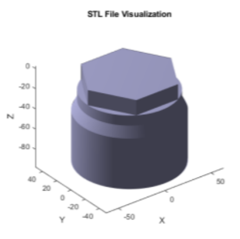	It is nearly a revolution part. Revolution axis: Z. Revolution grade: 99.23%.	The part is more suitable for machining in 3 axes.	No concave radius detected.	Detected 0 holes.	Yes

Part (STL view)	Output from algorithms				Accordance with expert
	Revolution symmetry	Geometric complexity	Minimum complex radius	Holes and classifications	
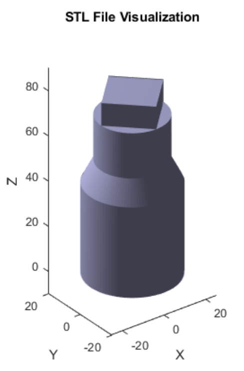	It is nearly a revolution part. Revolution axis: Z. Revolution grade: 99.16%.	The part is more suitable for machining in 3 axes.	No concave radius detected.	Detected 0 holes.	Yes
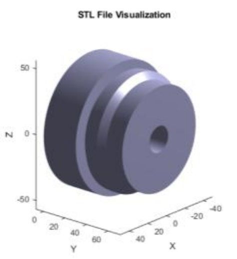	It is a revolution part. Revolution axis: Y.	The part is more suitable for machining in 3 axes.	The minimum concave radius is 10.00 mm	Detected 1 hole. 1 hole of radius = 10.00 mm and length = 40.00 mm.	Yes
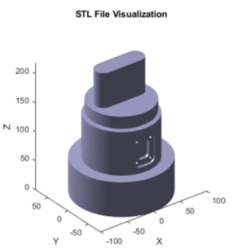	It is not a revolution part.	The part is more suitable for machining in 3 axes.	The minimum concave radius is 2.00 mm	Detected 6 holes. 6 holes of radius = 2.00 mm and length = 10.00 mm.	No. It could be classified as nearly revolution part
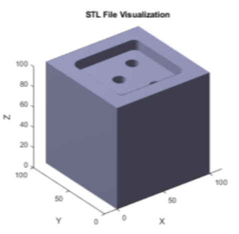	It is not a revolution part.	The part is more suitable for machining in 3 axes.	The minimum concave radius is 6.00 mm	Detected 3 holes. 3 holes of radius = 6.00 mm and length = 40.00 mm.	Yes
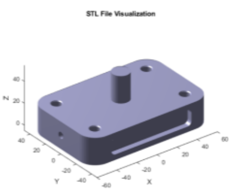	It is not a revolution part.	The part is more suitable for machining in 3 axes.	The minimum concave radius is 3.00 mm.	Detected 5 holes. 4 holes of radius = 4.25 mm and length = 19.25 mm. 1 hole of radius = 3.00 mm and length = 120.00 mm.	Yes
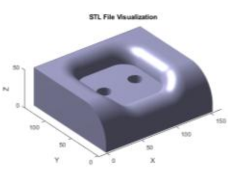	It is not a revolution part.	The part is more suitable for machining in 5 axes.	The minimum concave radius is 8.75 mm.	Detected 2 holes. 2 holes of radius = 8.75 mm and length = 30.00 mm.	Yes

Part (STL view)	Output from algorithms				Accordance with expert
	Revolution symmetry	Geometric complexity	Minimum complex radius	Holes and classifications	
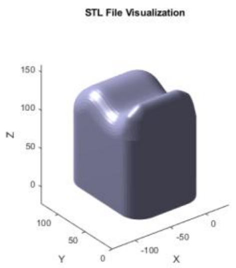	It is not a revolution part.	The part is more suitable for machining in 5 axes.	No concave radius detected.	Detected 0 holes.	Yes

In most cases, the algorithms produced results that were consistent with expert judgement. Only in one instance did the algorithm fail to correctly identify the number of holes, due to the singularity and high symmetry of a part that was rotationally symmetric. Nevertheless, the algorithm adopted a conservative approach in its estimation. In another case, the algorithm's classification of the part as not-revolution-part differed from the expert's view; however, the expert acknowledged the algorithm’s reasoning, noting that the assessment involves a degree of subjectivity.

Initial results for 2D drawing analysis

OCR based detection

A set of real 2D drawings for representative parts were evaluated. In some cases, 2D drawings include a table of tolerances located in a corner, rather than specifying them directly on the sketch. Basic OCR methods often face challenges when processing such layouts. While more advanced OCR approaches can detect these indicators, they may still encounter issues such as:

- ***Different font sizes.*** Annotations can be in different size, which can confuse some OCR analysis.
- ***Text alignment***. OCR can detect text aligned with different degrees, but it works better with 0º, 90º, 180º and 360º. The image can rotate to custom degrees, but that increases the time needed to analyse it.
- ***Bad quality***. Poorly scanned 2D sketches rely bad results.

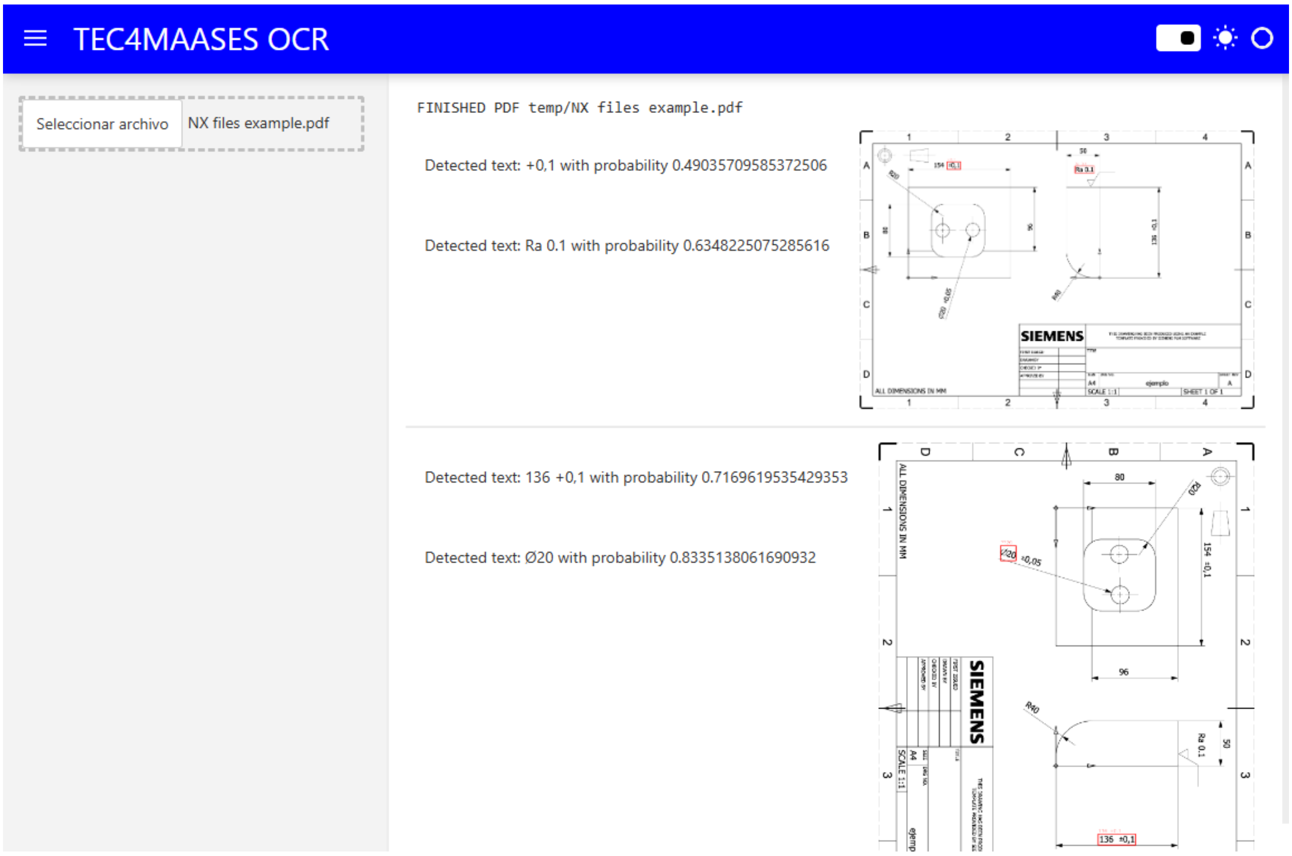

Figure 26: Example of 2D drawing analysis using OCR

Use of AI tools

While the initial results from using ChatGPT as a tool to detect surface roughness and dimensional tolerances from 2D drawing were promising in terms of accuracy, the risk of leaking confidential data put a stop on this path. The option of developing a local solution by integrating LLM and OCR libraries was evaluated but ultimately ruled out, as the development effort required would exceed the scope of the T4M project.

In addition, some exploratory activities were performed to evaluate the feasibility of using Large Language Models (LLMs) to detect features such as slots, holes, and steps from CAD models for improved process planning. In the study., ASCII-formatted STL files were treated as language-based data for a feature classification task. For doing so, Qwen, a pre-trained LLM with robust performance across diverse tasks, is compared to a CNN model that processes vowelised input. The results were validated to understand the capabilities and computational effort required by LLMs in the context of AFR, and to evaluate their understanding of new text-based formats such as STL. For this purpose, a dataset of 24 feature classes is used for a classification task. According to the results, LLM-based methods demonstrate an understanding of STL data, revealing potential for feature classification in this field. Results are reported in (Garcia et al. 2025).

3.2.5.2 Bill of Processes Generator: Processing time estimation results

To evaluate the accuracy of the processing time estimation tool, a comparative analysis was conducted between the T4M process time estimation tool and the theoretical machining time calculated using a set of CAM (Computer-Aided Manufacturing) programs. These CAM programs include the complete toolpaths and cutting conditions, meaning that the only discrepancies from actual production times are due to auxiliary operations such as part loading/unloading or tool changes. The evaluation began with a randomly selected sample of 25 milling CAM programs from Tekniker's internal database. An initial screening was then performed to exclude operations that did not meet the following criteria:

- The part must be manufactured from raw stock material (a metallic block) it cannot be a previously machined part.
- The raw stock must have a minimum machining clearance of 3 mm.
- The entire part must be machined, not just certain areas.

Before the screening, the estimates from the developed calculation tool were compared with the CAM predictions for all 25 programs. In general, significant deviations were observed.

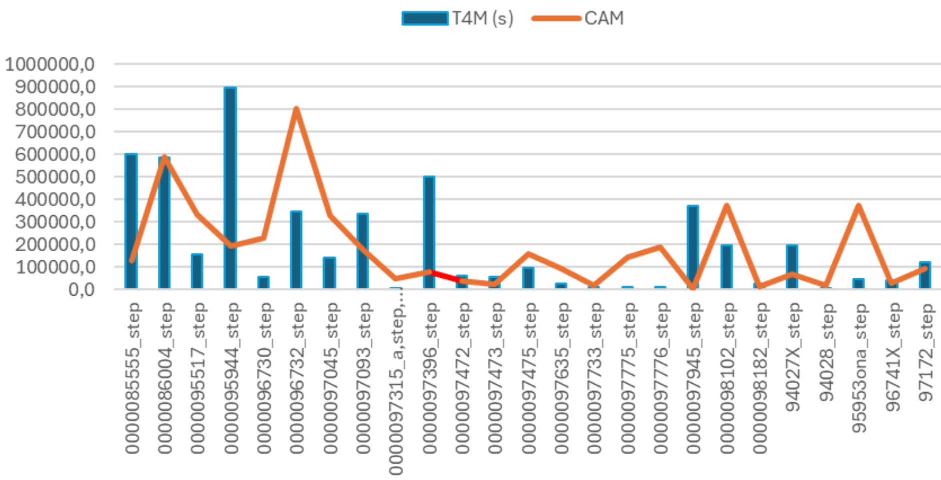

Figure 27: CAM estimation VS T4M process time estimation. 25 programs

After the screening process, it was concluded that only 7 parts were manufactured in accordance with the proposed manufacturing criteria. The main reason for the initial deviation was that many of the parts were not fully machined from raw stock but rather involved specific operations on prefabricated components. In such cases, the estimated time provided by the T4M tool is always higher, as it assumes the entire part must be machined from raw material. Another source of deviation was incorrectly declared cutting parameters. All parts were programmed assuming steel as the base material, but some CAM programs used cutting parameters of harder alloys. In these cases, the processing times estimated by the CAM software were significantly higher.

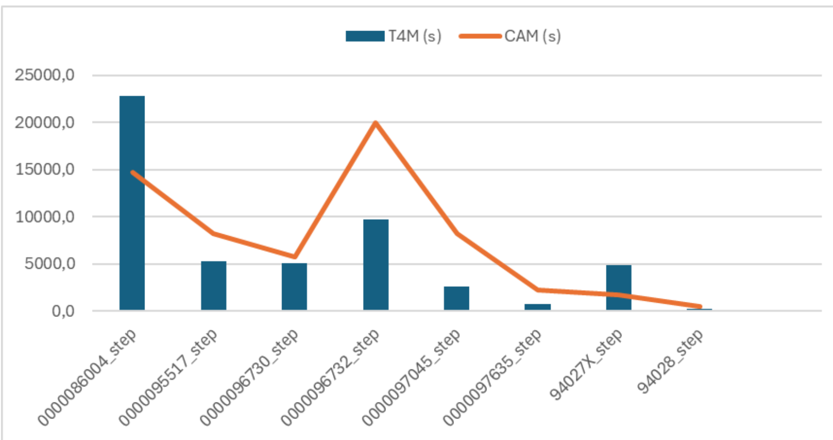

Figure 28: CAM estimation VS T4M process time estimation after screening. Only 7 parts satisfy the criteria.

The seven parts that meet the established criteria still show some degree of occasional deviation, but the overall trend is consistent. Among other aspects to improve, finishing operations play a significant role. In some cases, strict surface roughness requirements can lead to a higher number of finishing passes to achieve the desired surface quality.

3.2.5.3 Post-Optimization evaluation analysis

In addition to the development and validation of the Bill of Processes Generator, a post-optimization phase was conducted to assess the effectiveness and usability of the generated composition alternatives. This phase implements a two-step decision-support framework, incorporating *interactive filtering* and *multi-criteria ranking*, to aid users in identifying the most suitable supply chain configurations according to their preferences and priorities. Various test scenarios were conducted to observe how the resulting compositions change when using both filtering and ranking together, as opposed to using each method independently.

Table 13: Summary of the optimal compositions

ID	Number of Providers	Providers	Quality of Service Score
Composition 1	1	Tekniker	4.89
Composition 2	2	Tekniker, URA	4.00
Composition 3	1	URA	2.79
Composition 4	2	URA,Erreka	4.87
Composition 5	1	Erreka	2.18

In the first instance, synthetic data was used to demonstrate the interactive filtering and ranking process. Initially, five composition alternatives were generated, each defined by a specific provider set and quality-of-service score (see Table 1). The quality of service is represented by an index derived from historical data maintained by the platform, reflecting the past performance of each provider. The overall quality score for a composition is calculated as the average of the individual providers' quality scores within that composition. In future versions, this metric could be enhanced by incorporating a weighted average that takes into account user-defined preferences for each provider, offering a more personalized evaluation.

The user is able to assign preference scores to individual providers, where -1 indicates dislike, 0 is neutral, and 1 signifies a preferred provider. Additionally, the user can define two threshold constraints for the filtering process that are the maximum number of providers per composition and the minimum required quality score. These constraints help to deteriorate the set of feasible composition alternatives based on both qualitative and quantitative user-defined criteria. In this example we assume:

- Preferences: Erreka = 1, Tekniker = 1, URA = 1
- Maximum number of providers: 1
- Minimum quality score: 2.2

Based on these criteria, the filtering phase narrows down the alternative option to two viable compositions as shown in Table 14

Table 14 Remaining compositions after the filtering phase

ID	Number of Providers	Providers	Quality of Service Score
Composition 1	1	Tekniker	4.89
Composition 3	1	URA	2.79

Subsequently, the ranking phase is initiated. The user assigns weights to three evaluation criteria:

- Number of providers: 0.33
- Quality score : 0.33
- Preference score: 0.34

Using these weights, the system applies the TOPSIS (Technique for Order of Preference by Similarity to Ideal Solution) method to generate a final ranking table, enabling the user to identify the most suitable composition based on their specific priorities. The resulting rankings are presented in Table 15.

Table 15 Ranked compositions by TOPSIS

ID	Number of Providers	Providers	Quality of Service Score	TOPSIS score
Composition 1	1	Techniker	4.89	1.00
Composition 3	1	URA	2.79	0.50

Composition 1 emerged as the ideal configuration, demonstrating how the framework aligns results with user-defined priorities.

In the second scenario, the filtering phase was skipped. The same five compositions were considered, and the user provided neutral preferences for all the providers such as:

- Erreka: 0

- Tekniker: 0
- Ura: 0

The compositions and their scores are listed in Table 16.

Table 16: Composition for ranking without filtering

ID	Number of Providers	Providers	Quality of Service Score
Composition 1	1	Tekniker	2,65
Composition 2	2	Tekniker, URA	3,68
Composition 3	1	URA	2,34
Composition 4	2	Erreka, URA	5.00
Composition 5	1	Erreka	3,03

Applying the same weight distribution, the ranking results are presented in Table 17.

Table 17: Ranked compositions by TOPSIS (without filtering)

ID	Number of Providers	Providers	Quality of Service Score	TOPSIS score
Composition 5	1	Erreka	3.03	0.582
Composition 1	1	Techniker	2.65	0.533
Composition 3	1	URA	2.34	0.500
Composition 4	2	Erreka-URA	5.00	0.500
Composition 2	2	Techniker-URA	3.68	0.311

The absence of filtering resulted in a more dispersed range of scores, suggesting that non-ideal compositions were retained, potentially reducing decision clarity.

Finally, a comprehensive experiment was conducted using artificially generated compositions and provider configurations to examine the robustness of the decision-support framework under varying inputs. Multiple experiments were conducted by:

- Varying weights across the three criteria
- Assigning diverse user preferences (−1, 0, 1)
- Enabling/disabling filtering mechanism

Each run was recorded with an experiment ID and stored in structured .csv files. For the impact of filtering, an analysis comparing runs with and without filtering showed that maximum TOPSIS scores remained identical (1.000), indicating that the best solutions were always preserved. However, mean TOPSIS scores dropped slightly from 0.5377 to 0.5079 with filtering (see Table 18), suggesting a narrower but higher-quality solution space.

Table 18: TOPSIS experimentation analysis

Metric	Without Filtering	With Filtering
Mean Score	0.5377	0.5079
Max Score	1.000	1.000

A correlation analysis revealed a moderately positive relationship (r = 0.394) between user preference scores and final TOPSIS scores, affirming the system's ability to reflect user priorities without overshadowing other criteria. Further investigation (see Table 19) showed that high-quality compositions (Quality > 4) could still score highly in TOPSIS even when users gave neutral or negative preference values, reinforcing the model's balanced nature.

Table 19: High-Quality Compositions with Low Preferences

Experiment number	Composition	quality	score	providers
34	Composition 5	4.93	1.0	Erreka
39	Composition 5	5.0	1.0	Erreka
40	Composition 1	4.48	1.0	Tekniker
32	Composition 3	4.46	0.9038	URA
42	Composition 3	4.27	0.8964	URA
9	Composition 5	4.77	691	Erreka
38	Composition 4	4.44	0.6124	Erreka, URA
31	Composition 2	4.66	0.4782	Tekniker, URA
43	Composition 4	4.02	0.1813	Erreka, URA
92	Composition 1	4.97		Tekniker
93	Composition 5	4.98		Erreka
94	Composition 1	4.41		Tekniker
96	Composition 1	4.56		Tekniker

Notably, the results show that certain compositions with very low aggregate preference scores for their constituent providers can still achieve the maximum possible TOPSIS score (1.000). This suggests that the TOPSIS method effectively balances high-quality service scores with less favorable preference inputs, ensuring that top-performing compositions are not unjustly penalized when user preferences are not strongly aligned. Conversely, some compositions with high quality scores were excluded during the filtering phase and therefore do not appear in the final TOPSIS rankings. This demonstrates that neutral or negative user preferences alone do not automatically disqualify strong candidates, affirming the system's robustness and balance across multiple evaluation criteria. To further investigate the impact of filtering, the dataset was split into two subsets: one with filtering enabled and one without. For each subset, the Quality Score and Aggregate Preference Score were grouped into defined intervals. The average TOPSIS score was then calculated within each interval combination, resulting in two comparative heatmaps.

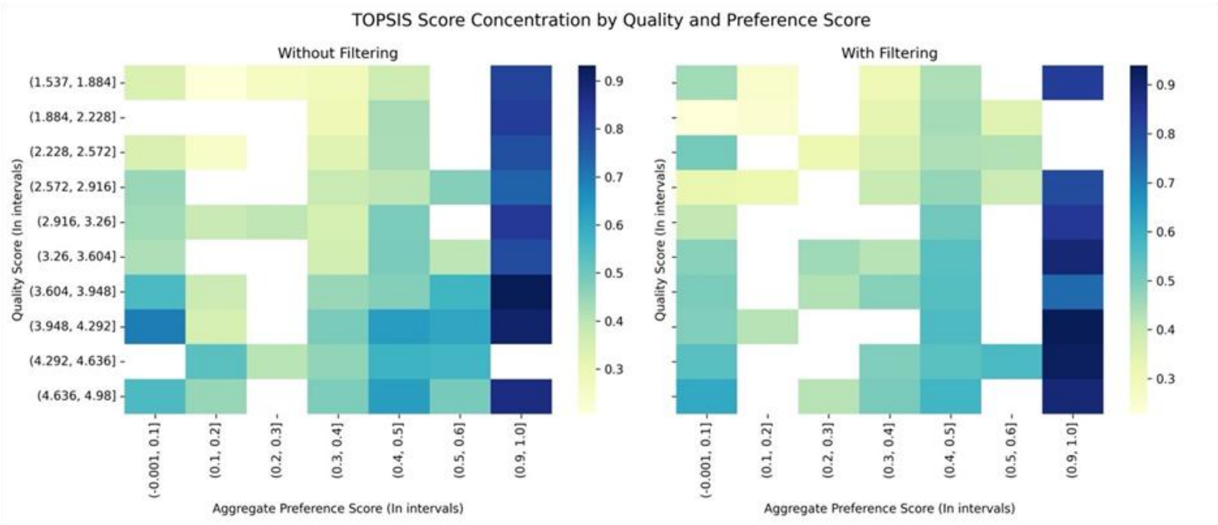

Figure 29: TOPSIS heatmap regarding Quality/Preferences

These visualizations reveal that without filtering, high TOPSIS scores are dispersed across a broader range of the heatmap. This means that even compositions with low quality scores and moderate preference scores can remain in the final list, potentially diluting the clarity of the results. In contrast, with filtering activated, the highest average TOPSIS scores become concentrated in the upper-right region of the heatmap, indicating that less competitive compositions were effectively removed. As a result, the final selection set becomes more focused and interpretable, simplifying the decision-making process for the user.

This two-phase decision-support approach mirrors established practices in the Multi-Criteria Decision Analysis (MCDA) literature:

A. Filtering is first applied to reduce the computational burden and eliminate unqualified alternatives (Lamrini et al., 2020)
B. TOPSIS is then employed as a robust ranking mechanism, commonly used in supply chain decision-making (Hwang & Yoon, 1981)

The three evaluation criteria—number of providers, quality, and user preference—are widely recognized in MCDA applications (Shahanaghi & Yazdian, 2009; Uygun & Dede, 2016; Azadeh et al., 2017). The inclusion of the filtering step acts as a safeguard, ensuring that weak or ambiguous configurations are removed early. This results in a more coherent and structured solution space across all decision dimensions.

Moreover, the modular nature of this approach supports traceability and adaptability—additional criteria can be incorporated or removed as needed without compromising the interpretability of the final results.

To further enhance usability and integration, the platform also supports data export in JSON format, allowing users to seamlessly transfer filtered and ranked configurations into external tools, share them with stakeholders, or incorporate them into broader analytics and decision-making workflows.

4 Conclusions and next steps

This deliverable presented the initial version of the analytics services supporting decomposition, resource matching, descriptive analytics and visualizations that encapsulate the composition service within the Tec4MaaSEs platform. Building on the requirements and architecture defined in WP2, this work introduces models and methods that enable dynamic configuration of service offerings, based on resource availability, service constraints, and performance objectives. Through pilot-specific developments in VN1 and VN2, the current implementation demonstrated how analytics driven pre and post processing can facilitate optimized service configurations and decision-making processes across diverse value networks. The developed modules ranging from decomposition tools and capability matching algorithms to interactive dashboards have laid a robust foundation for the realization of a scalable and resilient MaaS ecosystem. The next phase of development, to be detailed in Deliverable D3.4, will focus on the integration of these analytics components into a unified platform environment, enabling end to end execution of the service composition and decision support workflow. Special emphasis will be placed on validating system performance with larger (real and/or artificial) datasets and broader use case coverage across the pilot networks. This will allow for a more realistic evaluation of scalability and robustness in dynamic manufacturing scenarios.

Additionally, the visualization and dashboard services will undergo targeted enhancements. In particular, consumers will gain the ability to select their preferred composition alternative directly through the interface, initiating the negotiation phase with corresponding providers. This interaction marks a critical step toward decision making, where strategic preferences and trade-offs are seamlessly integrated into the service orchestration process.

As part of this enhancement, successful negotiation outcomes will be structurally decomposed into actionable tasks using an updated Decomposition Tree Schema. This schema will automatically allocate responsibilities across providers, facilitating the derivation of a coherent production schedule based on provider-specific capabilities, availability, and process constraints.

To summarize, further developments to be investigated and potentially materialized in D3.4 could include:

- Full integration of analytics modules with optimization and user interface components.
- Implementation of batching policies, as described in paragraph 2.1.2.3, and within the context of VN1.
- Execution of large-scale experiments for testing data pipeline performance, composition accuracy, and response times.
- Introduction of consumer-driven composition selection and negotiation functionalities.
- Deployment of a Decomposition Tree mechanism to operationalize agreed compositions into production schedules.
- Refinement of dashboard services with extended KPIs and real-time interaction capabilities.

These developments lay the foundation for the operational intelligence layer of Tec4MaaSEs, enabling the transition toward adaptive and context-aware Manufacturing-as-a-Service. In combination with the upcoming steps, they aim to transform the Tec4MaaSEs vision into a dynamic, explainable, and negotiation-ready MaaS platform that supports agile service composition and resilient manufacturing execution.

References

[1] Azadeh, A., Rahimi, Y., Zarrin, M., Ghaderi, A., & Shabanpour, N. (2017). A decision-making methodology for vendor selection problem with uncertain inputs. *Transportation Letters*, 9(3), 123-140.

[2] Azwir, H. H., & Christian, T. (2021). Optimizing Rapid Granulator Machines Utilization Using Overall Equipment Effectiveness and Capacity Measurement in Injection Molding. *Performa: Media Ilmiah Teknik Industri,* 20(2), 177–188.

[3] Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: forecasting and control (5th ed.). John Wiley & Sons.

[4] Bronja, H. (2011). MULTI-CRITERIA APPROACH TO RANKING SUPPLIERS IN THE SUPPLY CHAINS CONCEPT. Tehnicki vjesnik/Technical Gazette, 18(3).

[5] Garcia, O., de Calle, K. L., & Sarasua, J. A. (2025). Do LLMs understand shapes? Exploring STL files for automatic CAD feature recognition. *Procedia* CIRP, 133, 126-131.

[6] Gardner Jr, E. S., & McKenzie, E. (2011). Forecasting trends in time series. *Management Science,* 31(10), 1237-1246.

[7] Hayes, R. H., Wheelwright, S. C. (1979a). Link Manufacturing Process and Product Lifecycles. *Harvard Business Review*, no. January-February, 133-140.

[8] Hayes, R. H., Wheelwright, S. C. (1979b). The Dynamics of Process-Product Lifecycles. *Harvard Business Review*, no. March-April, 127-136.

[9] Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: principles and practice (2nd ed.). OTexts.

[10] Hopp, W. J., & Spearman, M. L. (2011). Factory physics (3rd ed.). Waveland Press.

[11] Hwang, C. L., & Yoon, K. P. (1981). Multiple attribute decision making: Methods and applications. Berlin/Heidelberg/New York: Springer-Verlag.

[12] Kutucuoglu, K., Hamali, J., Irani, Z., & Sharp, J. (2001). A framework for managing maintenance using performance measurement systems. *International Journal of Operations & Production Management*, 21(1/2), 173–195.

[13] Lamrini, L., Abounaima, M. C., Alaoui, M. T., El Mazouri, F. Z., El Makhfi, N., & Ouzarf, M. (2020, December). A filtering approach used in a massive data context to reduce the set of choices in a multicriteria decision aid process: Pareto solutions. In 2020 IEEE 2nd International Conference on Electronics, Control, Optimization and Computer Science (ICECOCS) (pp. 1-6). IEEE.

[14] Lu, Y., & Xu, X. (2017). A semantic web-based framework for service composition in a cloud manufacturing environment. *Journal of manufacturing systems*, 42, 69-81.

[15] Nahmias, S., & Olsen, T. L. (2015). Production and operations analysis (7th ed.). Waveland Press.

[16] Patel, H., & Shah, S. C. (2014). Review on cycle time reduction in manufacturing industries. J*ournal of Emerging Technologies and Innovative Research* (JETIR), 1(7), 955–957.

[17] Rimo, T. H. S., & Tin, O. C. (2017). A simulation study of capacity utilization to predict future capacity for manufacturing system sustainability. IOP Conference Series: Earth and Environmental Science, 109(1), 012020.

[18] Sahni, S., & Vagrecha, K. (2022). Reduction in Manufacturing Cycle Time cum Cost propels Productivity and Business growth in Aerospace Industries. *International Journal of Scientific and Research Publications* (IJSRP), 12, 340–351.

[19] Shahanaghi, K., & Yazdian, S. A. (2009). Vendor selection using a new fuzzy group TOPSIS approach. *Journal of Uncertain Systems*, 3(3), 221-231.

[20] Shahidul, M., Syed Shazali, S., Abdullah, Y., Ting, C., Hishamuddin, A., Azrin, M., & Adzlan, A. (2013). Measuring machinery capacity utilization and its impact on manufacturing performance and environment. *Journal of Manufacturing Operations Research and Sustainability*, July, 7–12.

[21] Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). Inventory and production management in supply chains (4th ed.). CRC Press.

[22] Stadtler, H., & Kilger, C. (Eds.). (2008). Supply chain management and advanced planning: concepts, models, software, and case studies (4th ed.). Springer.

[23] Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., & Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, 252(1), 1-26.

[24] Tao, F., Zhao, D., Hu, Y., & Zhou, Z. (2008). Resource service composition and its optimal-selection based on particle swarm optimization in manufacturing grid system. *IEEE Transactions on industrial informatics*, 4(4), 315-327.

[25] Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting*, 16(4), 437-450.

[26] Vollmann, T. E., Berry, W. L., Whybark, D. C., & Jacobs, F. R. (2005). Manufacturing planning and control for supply chain management (5th ed.). McGraw-Hill.

[27] Uygun, Ö., & Dede, A. (2016). Performance evaluation of green supply chain management using integrated fuzzy multi-criteria decision-making techniques. *Computers & Industrial Engineering*, 102, 502-511.

[28] Zilci, B. I., Slawik, M., & Küpper, A. (2015). Cloud service matchmaking using constraint programming. In 2015 IEEE 24th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (pp. 63-68). IEEE.