

Identification of Responsibility-Skill Coupling Patterns in Embodied Intelligent Manufacturing Driven by Improved BERTopic and Association Rules

Anyuan Zhong^{1*}

¹International Business School, Shandong Jiaotong University, 1508 Hexing Road, Shuangdao Bay Science and Technology City, Weihai, Shandong, China

Abstract

INTRODUCTION: Embodied Intelligence (EI) is rapidly penetrating industrial manufacturing, yet the interdisciplinary nature of EI systems makes it challenging for enterprises to precisely identify evolving talent demands.

OBJECTIVES: This study aims to identify responsibility topics and responsibility-skill coupling patterns in embodied intelligent manufacturing, thereby providing a theoretical basis for talent demand analysis.

METHODS: A total of 749 job advertisements were scraped and an improved BERTopic method integrating GSDMM topic similarity with semantic representation was proposed, identifying 14 responsibility topics (10 domain-specific) across four clusters. The Apriori algorithm was then applied to mine association rules between responsibility topics and skill keywords, with thresholds determined by the elbow method.

RESULTS: GSDMM-BERTopic consistently outperformed classic BERTopic and Gibbs-BERTopic, achieving average CV coherence of 0.5477, silhouette score of 0.6737, and outlier ratio of 0.1662. The selected configuration yielded 14 responsibility topics (10 domain-specific) across four clusters. Association rule mining produced 128 filtered rules. Notable topic-skill couplings include 3D Perception Modeling with Detection and Segmentation (lift 9.015), Motion Control with Reinforcement Learning and ROS (lift 2.277), and Frontier Technology Tracking with Ph.D. (lift 2.605). Cross-category analysis reveals Python and PyTorch as the common skill substrate, C++ as critical in hardware-adjacent domains, and the ROS/ROS2 ecosystem as the central middleware.

CONCLUSION: The findings reveal a T-shaped skill profile in embodied intelligent manufacturing and provide practical guidance for enterprises to optimize workforce allocation and talent cultivation strategies.

Keywords: Embodied Intelligence (EI), Smart Manufacturing, Job Advertisement Text Analysis, BERTopic, GSDMM, Association Rule Mining

Received on 10 August 2026, accepted on 14 September 2026, published on 02 October 2026

Copyright © 2026 Anyuan Zhong., licensed to EAI. This is an open access article distributed under the terms of the [CC BY-NC-SA 4.0](#), which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

doi: 10.4108/eetsis.14410

1. Introduction

Embodied Intelligence (EI), as the latest paradigm in artificial intelligence development, is rapidly penetrating into industrial manufacturing and has become a key enabling technology driving the transformation and upgrading of smart manufacturing [1]. Driven by groundbreaking advances in the new generation of artificial

intelligence, the Embodied Intelligence-Empowered Smart Manufacturing (ESME) ecosystem is breaking down the boundaries among heterogeneous elements inherent in conventional manufacturing. Through dynamic closed-loop synergy across three interconnected stages, namely design optimization, manufacturing execution, and

*Corresponding author. Email: 223013@sdjtu.edu.cn

operational evolution, industrial robots are being endowed with enhanced capabilities in perception, understanding, and complex task planning and execution [2]. With the continuous progress of embodied intelligence, large language models, and cloud computing, Industry 5.0 promotes industrial artificial intelligence development through human-centered cyber-physical-social systems. Large language models, via breakthroughs in unimodal capabilities, enhancement through multimodal fusion, and embodied realization, enable industrial robots to possess stronger perception, comprehension, and complex task planning and execution abilities [1, 3].

However, the development and deployment of embodied intelligent manufacturing systems are typical interdisciplinary system engineering endeavors, requiring the integration of multidisciplinary competencies including algorithm development, control systems, mechanical design, communication architecture, and system integration. This makes talent demands characterized by a wide span and rapid changes. Different manufacturing scenarios exhibit significant variations in skill demand patterns: for example, assembly scenarios emphasize fine manipulation and collaborative capabilities, inspection scenarios prioritize perception and data analysis capabilities, while flexible production scenarios demand rapid adaptation and reconfiguration capabilities. The manufacturing sector is currently facing a severe skills gap, prominently manifested as "limited transferable skills" and "difficulties in transitioning from content knowledge to practical competencies" [4]. There is an urgent need for interdisciplinary talents capable of bridging multidisciplinary knowledge and rapidly adapting to technological iterations, yet a significant disconnect exists between current talent supply and the integration of theory and practice. This misalignment of talent capabilities directly leads to efficiency losses and increased costs for enterprises in new technology deployment, production line adaptation, and anomaly response [4].

The numerous online job advertisements are a type of unstructured data that reflects the current technological needs of enterprises and provides relatively timely information about talent requirements in the labor market [5]. A considerable number of studies have applied text mining techniques to extract skill information from job advertisements, such as topic modeling, keyword extraction and rule-based methods [5-8]. A limitation of prior studies is that they only focus on which skills are needed for a specific job and do not explore the relationship between these skills and job responsibilities. Skills are often associated with particular dimensions of responsibility, and these responsibility areas have different patterns of skill requirements. Without analysis of the coupling between responsibility and skills, it becomes difficult to know exactly how to address changes in skill requirements in different situations for manufacturing enterprises.

To solve this problem, this study collected 749 job advertisements related to embodied intelligent manufacturing and proposed an improved BERTopic

(Bidirectional Encoder Representations from Transformers Topic) method [9] that integrates GSDMM (Gibbs Sampling Dirichlet Multinomial Mixture) [10] and semantic representation for topic modeling of job responsibilities to find different responsibility topics. Next, association rule mining was applied to identify associations between responsibility topics and skill keywords, based on which a fine-grained competency requirements model was constructed. The model can reveal the pattern of responsibility-skill coupling in all kinds of manufacturing conditions. Building upon this framework, the model offers a theoretical foundation and practical reference for enterprises to understand shifts in talent demand, and helps them optimize the allocation and training of talent. The three main contributions of this paper are as follows: (1) To propose an improved BERTopic method that combines GSDMM and semantic representation to increase the accuracy and interpretability of responsibility topic modeling; (2) To build a framework for discovering responsibility-skill association rules and reveal the characteristics of skill demand in various manufacturing scenarios; (3) To verify the practical value of the proposed method on real job advertisement data and provide a technical path for talent demand analysis in embodied intelligent manufacturing.

2. Related Works

Embodied Intelligence (EI), as a transformative paradigm in the field of artificial intelligence, represents the deep integration of multimodal perception, cognitive reasoning, and physical environment interaction, and is increasingly becoming a key enabling technology for the development of smart manufacturing [11]. Existing manufacturing systems generally rely on offline trained models or predefined rules, embodying a "disembodied cognition" pattern. When faced with environmental disturbances, task reconfiguration, or unstructured scenarios, they suffer from decision-making delays or behavioral rigidity. In contrast, embodied intelligence, by constructing a "perception-decision-execution" closed loop, provides a technical framework and evolutionary pathway for achieving autonomous and agile smart manufacturing [11]. In the domain of robotic manipulation, physics-informed embodied intelligence frameworks integrate data-driven learning with physical laws and constraints, offering new approaches to solving object spatial reconstruction problems under complex environmental constraints. Particularly in modern manufacturing processes that require robots to consider force, morphology, and even material properties, breakthroughs in multiphysics modeling capabilities are urgently called for [12]. As Industry 5.0 transitions from efficiency-oriented automation to human-centric, resilient, and sustainable manufacturing paradigms, human-robot collaboration has become a key development direction. Generative foundation models provide enabling mechanisms for human-robot collaboration across three levels—

perception, cognition, and execution—driving collaborative modes from predefined collaboration toward adaptive collaborative intelligence [13]. These studies collectively demonstrate that embodied intelligence is not only a frontier direction of technological development but also a core driving force for realizing the transformation and upgrading of smart manufacturing.

Job advertisements are a type of data that reveal the types of employees that enterprises seek, making them a widely used data source in research. Researchers have used text mining, content analysis and other methods to extract skill requirement information from them [14, 15]. For example, Pejic-Bach et al. have conducted text mining on job advertisements related to Industry 4.0 and found that there are two main categories of job characteristics: one focuses on core Industry 4.0 knowledge, and the other covers general knowledge domains such as supply chain management and enterprise software [14]. Verma et al. have carried out content analysis of job advertisements for artificial intelligence and machine learning positions and found that machine learning jobs have placed a relatively high demand on technical skills (such as data mining, programming and statistics), whereas artificial intelligence jobs require more soft skills, such as communication [16]. Wang et al. proposed the JobViz system to help job seekers learn about the necessary skills visually [17]. Beyond skill extraction, existing studies have also examined the impact of gendered language in job ads on applicant behavior [18] and proposed text-level debiasing methods for job advertisements [19]. However, there is one existing problem in the field: most studies have focused only on "what skills the job requires" and have not explored the underlying links between job responsibilities and skills. Gurcan et al. used BERTopic-based neural topic modeling to analyze job responsibilities and map skill categories to expertise roles[6], but this mapping remained descriptive and did not employ association rule mining to quantitatively uncover coupling patterns between responsibilities and specific skill combinations. Zhong et al. have suggested that job tasks, as a link between positions and skills, have not been paid much attention to. They put forward the "position-task-skill" demand model and used word co-occurrence network clustering to explore the intrinsic relationships among these three elements. However, this model is limited in showing the links among skill demands in different responsibility areas clearly [20]. Because there is no coupling between mining responsibility and skill, it is difficult to link the existing analysis of skill demand with specific working conditions; therefore, changes in skill demand in different manufacturing environments cannot be understood well, and this information is not practically useful for enterprises' talent allocation and development decisions. Association rule mining is a well-established technique for discovering hidden relationships among data entries. However, existing studies have not fully leveraged the structured responsibility information derived from topic modeling to guide association rule generation. Consequently, the mined

rules tend to lack semantic interpretability and domain specificity.

Responsibility descriptions in job advertisements are typically very brief, exhibiting typical short-text characteristics such as sparsity and high dimensionality, which pose challenges to traditional topic modeling methods [10]. Researchers have proposed various solutions for short-text clustering, among which GSDMM (Gibbs Sampling Dirichlet Multinomial Mixture) [10] and BERTopic (Bidirectional Encoder Representations from Transformers Topic) [9] are two representative approaches. GSDMM is a probabilistic model-based short-text clustering method that employs a Dirichlet Multinomial Mixture model and automatically infers the number of clusters through collapsed Gibbs sampling. It can effectively handle the sparsity and high-dimensionality issues of short texts and obtain representative words for each cluster[10]. However, GSDMM models text based solely on lexical (term) co-occurrence statistics, ignoring semantic information, making it difficult to capture deep semantic associations between words. It also lacks the ability to identify noisy data. In contrast, BERTopic is a neural topic modeling method based on semantic vectors. It generates document embeddings using pre-trained Transformer models, clusters these embeddings, and finally generates topic representations through class-based Term Frequency-Inverse Document Frequency (c-TF-IDF) [9]. BERTopic can effectively capture the semantic information of texts and generate topics with good coherence. However, its performance is less stable when applied to domain-specific texts, especially on small-scale datasets. To address these shortcomings, some studies have proposed incorporating word probability information into the BERTopic framework. For instance, Zhu et al. proposed a Gibbs-BERTopic hybrid method, which iteratively optimizes topic distributions through Gibbs sampling within the BERTopic framework, significantly improving the robustness and accuracy of the topic model and effectively reducing the impact of noise and data sparsity on topic extraction [21], but this approach refines the clustering results post hoc, meaning that if the initial clustering is biased, the subsequent smoothing can only correct within that biased framework rather than incorporating statistical information from the outset.

In summary, existing research on talent demand analysis in embodied intelligent manufacturing suffers from the following deficiencies: First, skill demand analysis mostly remains at the level of skill lists, lacking in-depth mining of responsibility-skill coupling relationships. This makes it difficult to explain the rationale behind skill demands and hinders the understanding of skill demand variations across different manufacturing scenarios. Second, the integration of association rule mining and topic modeling is insufficiently deep, failing to fully leverage structured responsibility topic information. Finally, although progress has been made in short-text topic modeling methods, GSDMM neglects semantic information, BERTopic is susceptible to noise in specific domains, and existing hybrid methods have not yet fully integrated word

probability information with semantic representation. To address these issues, this paper proposes an improved BERTopic method integrating GSDMM and semantic representation. It comprehensively utilizes the keyword co-occurrence perspective provided by GSDMM and the semantic vectors from BERTopic to compute the similarity of responsibility statements, and designs a customized BERTopic based on this similarity. Then, association rule mining is employed to find association rules among the responsibility topics and skill keywords, and thus a more accurate model of competency demand is constructed.

3. Methodology

This chapter presents the overall methodology framework, which is briefly shown in Figure 1. Firstly, job responsibilities and skill requirements were acquired through web crawling, filtering, and keyword computation. Secondly, given the characteristics of short texts, this study proposed an improved BERTopic algorithm that combines GSDMM and semantic embeddings, and performed topic modeling on job responsibilities. Lastly, association rules between responsibility topics and skill keywords were mined using the Apriori algorithm.




Figure 1. The Methodology Framework

3.1. Data Acquisition and Preprocessing

The three steps of data collection and preprocessing in this paper are: (1) data crawling; (2) data filtering and information extraction; and (3) word segmentation and keyword computation.

Playwright was employed to simulate a browser and obtain data from Baidu Baipin, a comprehensive recruitment search engine that aggregates job postings from multiple platforms. Playwright is relatively easy to set up, and it is faster than Selenium in headless mode [22]. To ensure that the collected posts are highly relevant to the field of "Embodied Intelligence", only that keyword was

used in the search. A job search was performed with the above keyword, and then, one by one, each of the detail pages was visited to obtain its raw HTML content. BeautifulSoup was then employed to parse the HTML and extract the structured data of job descriptions and job titles. The first crawl found about 1,200 job openings, and after removing duplicates, more than 1,000 were finally obtained.

Qwen3.7-max was used in this paper to filter the crawled data and extract information at the data filtering stage. To reduce the problem of bias in a single model, two stages of verification were employed; first, Qwen3.7-max[23] extracted the information, and then DeepSeek-

V3[24] was used to cross-verify by comparing the extracted entries with the original text. During extraction, the model filtered the list to include only technical positions in the field of embodied intelligent manufacturing and then extracted the job responsibilities and skill requirements from the job descriptions. According to the general rules of prompt engineering, the prompt was constructed with roles, rules and a few-shot example, and then improved in a pilot study on 20 job postings. As per the instructions in the prompt, the model was to extract the original text from the job descriptions for responsibilities and skills without any modifications, additions or deletions. All the responsibility items were required to be atomic; if one responsibility contained several tasks, these also needed to be listed separately. If the position was not a technical position in the field of embodied intelligence, an empty list should be returned. The output of the model was in JSON format, and the two keys were "responsibilities" and "requirements". After LLM-based filtering, a total of 749 valid job postings were obtained, with 4,973 responsibility items and 4,816 skill requirement items. To guarantee the quality of the data, 100 job postings were selected randomly and manually verified in this study. The manual verification confirmed an extraction accuracy of 0.95.

At the stage of word segmentation and keyword computation, the Jieba Chinese word segmentation library was used to tokenize the text. To improve the accuracy of segmentation, new words specific to the domain were extracted before segmentation based on information entropy and mutual information, and these identified domain-specific terms were added to the segmentation dictionary. At the same time, some useless stop words were added to the stop word list and excluded during segmentation. Word segmentation was carried out separately on the responsibility section and the skill requirements section of each job posting to obtain keywords for the next steps of topic modeling and association rule mining.

The preprocessing pipeline above produced a dataset of 749 job postings, 4,973 responsibility items and 4,816 skill requirement items, which provide a good data foundation for the topic modeling of responsibilities and skill-association rule mining in this paper.

3.2. Improved BERTopic

Compared to classic BERTopic, which relies solely on semantic embedding similarity for UMAP dimensionality reduction, the proposed method also incorporates GSDMM-based topic co-occurrence similarity, capturing short-text word co-occurrence patterns that pure semantic embeddings may overlook. Compared to Gibbs-BERTopic, which applies Gibbs sampling as a post-hoc smoothing step after the initial clustering, the proposed method integrates GSDMM similarity directly into the UMAP similarity matrix, allowing statistical and semantic signals to guide

dimensionality reduction simultaneously rather than smoothing potentially biased results after the fact.

3.2.1. Computation of GSDMM Topic Similarity

The hyperparameters of GSDMM were determined through a grid search. For short text clustering, both α and β are typically set to small values around 0.1, as a smaller α encourages sparser document-topic assignments, which tends to reduce the number of active topics, while a smaller β reduces the smoothing effect of the Dirichlet prior, allowing the observed word frequencies to dominate the topic-word distributions. The search ranges were set to $\alpha \in \{0.05, 0.1, 0.15, 0.2\}$ and $\beta \in \{0.05, 0.1, 0.15, 0.2\}$, centered around the conventional value of 0.1. Three independent repetitions were carried out for each set of parameters to obtain the mean coherence and perplexity. A weighted sum of normalised coherence (weight 0.6) and perplexity (weight 0.4) was used to find the best parameters. Given the intention for business applications, a relatively high weight is given to coherence, and thus semantically coherent topics can be employed to label occupational scenarios. A weight of 0.4 is given to perplexity to prevent the model from being too focused on semantic coherence and losing information from the data; otherwise, it would result in excessively general or redundant topics. The Parameter Search Space is:

$$\begin{aligned}\alpha &\in \{0.05, 0.1, 0.15, 0.2\} \\ \beta &\in \{0.05, 0.1, 0.15, 0.2\}\end{aligned}\quad (1)$$

The best values for alpha and beta were found to be 0.05 and 0.15. Given the above parameters, the coherence (c_v) was 0.5824 ± 0.0033 , the perplexity was 13.0977 ± 0.2401 , and the number of topics was 112. Although both coherence and perplexity showed satisfactory results, the number of topics (112) was too large for direct analysis, so BERTopic was used to reduce the number of topics.

For the responsibility corpus, topic similarity was computed iteratively with the number of topics starting from 5 and increasing in steps of 5. The underlying rationale is that if two responsibility items belong to the same topic within a given topic model, they are considered to share a certain degree of similarity, and the similarity magnitude is inversely related with the number of topics. For example, if responsibility items a and b belong to the same topic under a model with 5 topics, the similarity contribution is smaller than under a model with 50 topics.

Specifically, if a and b are assigned to the same topic in the topic model (with topic number n), the similarity is computed using the following formula:

$$\text{sim}_n(a, b) = \log_2(n) \cdot \sqrt{\text{score}(a) \cdot \text{score}(b)} \quad (2)$$

where $\text{score}(a)$ and $\text{score}(b)$ are the probabilistic scores of responsibility items a and b within that same topic. If a and b are not in the same topic, then $\text{sim}_n(a, b) = 0$. The GSDMM topic similarity is obtained by summing the similarities across all topic numbers:

$$sim_raw_{gsdmm}(a,b) = \sum_{n=5}^{100} sim_n(a,b)$$

$$sim_{gsdmm}(a,b) = 0.5 + 0.5 * \frac{sim_{gsdmm}(a,b)}{max(sim_{gsdmm})} \quad (3)$$

The normalization to [0.5, 1.0] rather than [0, 1] reflects a design consideration: if two responsibility items have ever been co-assigned to the same topic at any topic number n , their co-occurrence carries inherent semantic meaning and should contribute a non-negligible baseline similarity.

3.2.2. Customized BERTopic

BERTopic is a modular topic clustering model consisting of three main steps: dimensionality reduction of semantic vectors using Uniform Manifold Approximation and Projection (UMAP), clustering of the reduced vectors using Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN), and topic keyword generation using class-based Term Frequency-Inverse Document Frequency (c-TF-IDF).




Figure 2. Performance Comparisons of BERTopic, Gibbs-BERTopic, and GSDMM-BERTopic

The customization in this paper is reflected in the dimensionality reduction step: dimensionality reduction can be performed based on pre-computed similarity, where the similarity in this paper is the maximum of the GSDMM topic similarity and the semantic vector similarity:

$$similarity = \max(similarity_{gsdmm}, similarity_{semantic}) \quad (4)$$

where $gsdmm_similarity$ denotes the GSDMM topic similarity and $semantic_similarity$ denotes the semantic vector similarity. The semantic vectors were generated using the Qwen3 Embedding model[25].

After computing the similarity, BERTopic employed HDBSCAN for clustering to obtain topics. The entire BERTopic process involves two key hyperparameters: $n_neighbors$ in the UMAP dimensionality reduction step and $min_cluster_size$ in the HDBSCAN clustering step.

All three algorithms, namely classic BERTopic, Gibbs-BERTopic, and the proposed GSDMM-BERTopic, were evaluated across $n_neighbors \in \{4, 8, 16, 20, 24, 28, 32\}$ and $min_cluster_size \in \{20, 30, 40, 50, 60, 70, 80, 90, 100\}$, and compared on CV coherence, silhouette score, and outlier ratio. The full comparison results are visualized as heatmaps in Figure 2. As is shown in the Figure, GSDMM-BERTopic consistently outperformed the other two algorithms across the parameter grid. On average, GSDMM-BERTopic achieved a CV coherence of 0.5477 and a silhouette score of 0.6737, compared to 0.4027 and 0.3909 for classic BERTopic, and 0.4286 and 0.5082 for Gibbs-BERTopic. The outlier ratio of GSDMM-BERTopic (0.1662) was also lower than that of Gibbs-BERTopic (0.2919) and comparable to classic BERTopic (0.1808). Classic BERTopic produced an average of only 3.6 topics, which is insufficient for capturing the diversity of occupational scenarios, while Gibbs-BERTopic generated an excessive 29.5 topics on average with lower coherence. GSDMM-BERTopic struck a better balance, yielding an average of 18.1 topics with both higher coherence and higher silhouette scores.

3.3. Responsibility-Skill Association Rule Mining

In this framework, responsibilities are organised into categories according to different job functions, and skills are kept as individual keywords to preserve their original technical meaning. This design is based on the premise that responsibilities are narrative descriptions of job functions that can benefit from semantic grouping, whereas skills are individual technical terms that are best represented at the keyword level. By mining association rules between responsibility topics and skill keywords, the framework reveals not just "what skills are needed" but "what skill combinations are required for each type of job function," an insight that cannot be obtained from keyword-only association mining.

3.3.1. Association Rule Computation

Each job posting corresponds to one transaction. A transaction contains the responsibility topics and skill requirement keywords involved in that job posting. The responsibility topics were obtained through improved BERTopic clustering, and the skill requirement keywords were extracted from the skill requirements section of the job advertisement. The format of a transaction is:

$$\{Topic_1, Topic_2, \dots, Topic_k, Skill_1, Skill_2, \dots, Skill_m\} \quad (5)$$

where $Topic_i$ denotes the responsibility topic label and $Skill_j$ denotes a skill requirement keyword.

After generating the transactions, the Apriori algorithm[26] was employed to mine association rules between responsibility topics and skill requirement keywords. The association rules were evaluated using three metrics: Support, Confidence, and Lift. The form of an association rule is:

$$\{Topic_i\} \rightarrow \{Skills\} \quad (6)$$

The formulas for Support, Confidence, and Lift are as follows:

$$\begin{aligned} Support(A \rightarrow B) &= \frac{|A \cap B|}{|T|} \\ Confidence(A \rightarrow B) &= \frac{|A \cap B|}{|A|} \\ Lift(A \rightarrow B) &= \frac{Support(A \rightarrow B)}{Support(A) \times Support(B)} \end{aligned} \quad (7)$$

T is the total number of transactions, A is a responsibility topic, and B is a set of skill requirement keywords.

3.3.2. Association Rule Filtering

After completing the GSDMM-BERTopic topic modeling, a single parameter configuration was selected for rule filtering and subsequent analysis. The selection principle balanced three criteria: high coherence and silhouette score, a low outlier ratio, and a manageable number of topics to avoid overfitting and excessive analytical complexity. The configuration with $n_neighbors = 32$ and $min_cluster_size = 70$ (first run) was chosen, yielding a coherence of 0.63, a silhouette score of 0.71, an outlier ratio of 0.23, and 14 topics. The rationale is grounded in the elbow method: when metric values (confidence or support) are sorted in descending order and plotted against their rank, the values initially decline sharply, then flatten into a long tail of uniformly small and uninformative values. The elbow point marks the transition between these two regimes and serves as a natural threshold, balancing the retention of meaningful rules with the removal of noise. Confidence was treated globally because it represents a conditional probability, i.e., the likelihood of a skill combination occurring given a responsibility topic, which is comparable across different topics. Support, in contrast, measured the absolute frequency of a rule in the dataset, and since

different topics contain different numbers of documents, topic-specific thresholds are necessary.

The filtering procedure consists of five steps:

Step 1: Initial rule generation. Apriori's algorithm was employed to find association rules with the following relaxed thresholds: minimum support = 0.01, minimum confidence = 0.1 and minimum lift = 1.0. The maximum length of the rule is 3 (that is, a combination of up to three skill keywords). The lower threshold was not restrictive, thereby avoiding exclusion of potentially valuable rules.

Step 2: Global confidence threshold. All the candidate rules were sorted by confidence in descending order. Build the confidence-versus-rank curve, and then, with the help of the Kneedle algorithm, found the elbow point that represents the maximum curvature of the normalised curve. The detected elbow point has a global confidence threshold of 0.2062 and a curvature of 0.69; it is a relatively clear elbow (Figure 3). Rules with confidence below this threshold were considered too weak to be informative and are thus disregarded.




Figure 3. Confidence elbow (global)

Step 3: Topic-specific support thresholds. For each responsibility topic, the rules belonging to that topic were sorted by support in descending order, and the elbow point was detected on the support-versus-rank curve using the same Kneedle algorithm[27]. Since different topics have different numbers of associated documents, the support thresholds are topic-specific, ranging from 0.0170 to 0.1132. The curvature values at the detected elbow points range from 0.45 to 0.69, confirming that all elbow points are well-defined (Table 1). Rules with support below their topic-specific threshold are discarded.

Table 1. Thresholds and Kneedle Curvature

Parameter	Threshold	Kneedle Curvature
confidence (global)	0.206	0.695
support (topic_0)	0.0311	0.676

support (topic_1)	0.0212	0.656
support (topic_2)	0.0453	0.623
support (topic_3)	0.017	0.683
support (topic_4)	0.041	0.647
support (topic_5)	0.0269	0.673
support (topic_6)	0.113	0.539
support (topic_7)	0.0453	0.553
support (topic_8)	0.0509	0.647
support (topic_9)	0.0537	0.666
support (topic_10)	0.0354	0.455
support (topic_11)	0.024	0.642
support (topic_12)	0.017	0.492
support (topic_13)	0.0269	0.614

Step 4: Combined filtering. Rules that pass both the global confidence threshold and the topic-specific support threshold were retained. The two-level filtering will only keep the rules that have a high degree of statistical reliability and are also relatively frequent in their respective topics.

Step 5: Top-20% selection. Within each topic and each rule length (1, 2, or 3), the remaining rules are ranked by confidence, and the top 20% are selected. This step further focuses the analysis on the most salient rules while keeping the number of rules manageable for interpretation.

4. Results and Analysis

4.1. Responsibility Topic Analysis

In this study, the improved BERTopic method was applied to model the responsibility descriptions of embodied intelligent manufacturing positions. A total of 14 responsibility topics were identified, of which 10 pertain to specific embodied intelligent manufacturing scenarios. The remaining 4 topics are general-purpose descriptions (e.g., team collaboration, project management, technical documentation, and cross-departmental communication). These are excluded from detailed analysis as they are transversal competencies applicable across all scenarios rather than specific to embodied intelligence. Nevertheless, their emergence as distinct clusters suggests that collaborative and organizational capabilities are valued as independent competency dimensions alongside technical expertise. Based on their intrinsic associations, the 10 responsibility topics were further categorized into four clusters: Perception and Data Engineering, Algorithm and Model Development, Simulation, Control and Deployment, and System Integration and Testing. The keywords and the number of responsibility entries for each topic are shown in Table 2 below. Based on domain knowledge, it is reasonable to speculate that these categories exhibit natural affinities with typical manufacturing contexts: Perception and Data Engineering is closely tied to inspection tasks,

Simulation, Control and Deployment is essential for assembly operations, and Algorithm and Model

Development underpins flexible production that requires adaptive intelligence.

Table 2. Responsibility Topics

Topic	Keywords	Number
Testing, Integration and Commissioning	Communication, Joint debugging, Documentation, Process, Debugging, Testing, Equipment, Linux, Motor, Test case	855
Multimodal Data Governance	Annotation, Cleaning, Quality inspection, Data quality, Video, Version, Quality, Storage, Automation, Data collection	467
Motion Control	Grasping, Motion control, Obstacle avoidance, Decision-making, Dynamics, Trajectory, Joint, Kinematics, Machinery, Control	347
Data Loop and Evaluation	Evaluation, Simulation, Regression, Evolution, Training, Agent, Real robot, Quantization, Data loop, Model training	335
Multimodal Embodied Modeling	Language, Vla, World model, Vlm, Vision, Pretraining, Multimodal, Large model, Imitation learning, World	279
Frontier Technology Tracking	Frontier, Paper, Influence, Top conference, Frontier technology, Cvpr, Innovation, Icra, Achievement, Progress	218
Simulation Loop Construction	Isaac, Sim, Mujoco, Simulation, Lab, Simulation environment, Rendering, Gazebo, High fidelity, Sim2real	159
Model Deployment and Inference Optimization	Inference, Acceleration, TensorRT, Quantization, Computation, Deployment, Edge, Gpu, Device-side, High performance	134
Multi-Sensor Integration and Fusion	Imu, Sensor, Calibration, Synchronization, Rgb, Camera, Joint, Actuator, Fusion, High precision	84
3D Perception Modeling	Estimation, Detection, 3d, Vision, Reconstruction, Point cloud, Segmentation, Pose, Semantics, Posture	75

4.1.1. Perception and Data Engineering

This category comprises three responsibility topics: Multimodal Data Governance (467 entries), Multi-Sensor Integration and Fusion (84 entries), and 3D Perception Modeling (75 entries), totaling 626 entries. These three topics constitute a complete perception pipeline from data acquisition and sensor fusion to 3D perception.

Multimodal Data Governance covers the full lifecycle of training data management, including data collection, annotation, cleaning, quality inspection, and version control. With 467 entries, it is the largest topic within this category, reflecting that the growing demand for large-scale, high-quality training data in embodied intelligence has made data engineering a critical and specialized responsibility. The emphasis on data quality and versioning suggests that enterprises are moving from ad-hoc data practices toward systematic data governance.

Multi-Sensor Integration and Fusion refer to the calibration, time synchronization and data fusion of all sorts of sensors, such as IMUs, cameras and RGB sensors. Although there are only 84 entries, this topic is very specialized and serves as the hardware-software interface for perception systems.

The parts of 3D Perception Modeling are 3D Object Detection, Point Cloud Reconstruction, Semantic Segmentation and Pose Estimation. With only 75 entries, this topic is the smallest among all 10 topics but requires the highest level of algorithmic expertise and is concentrated in positions for algorithm R&D.

The three topics together form the "sensory" foundation of embodied intelligent systems and cover the whole

process of raw data governance and high-level perception algorithms.

4.1.2. Algorithm and Model Development

The three topics in this category are Multimodal Embodied Modeling (279 entries), Data Loop and Evaluation (335 entries), and Frontier Technology Tracking (218 entries), with a total of 832 entries; thus, algorithmic innovation is still the main driver of embodied intelligence.

Multimodal Embodied Modeling is one of the main algorithmic directions, including the development of VLA (Vision-Language-Action) models, world model construction, visual language modeling and imitation learning. The emergence of keywords such as "VLA", "world model", "VLM" and "pretraining" shows that the area of end-to-end multimodal systems for perception, reasoning and action is developing.

Data Loop and Evaluation addresses the engineering challenge of systematically assessing and improving embodied intelligent models. This includes model evaluation in simulation environments, automated regression testing, data loop iteration, and Agent-based training optimization. The 335 entries suggest that the industry has progressed from the initial stage of "can it work?" to "can it work stably and efficiently?" and that systematic evaluation workflows are now essential.

Frontier Technology Tracking involves monitoring top-tier conference publications (CVPR, ICRA, NeurIPS), tracking technological innovations, and translating academic advances into industrial applications. With 218 entries, this topic reflects the high sensitivity of the

embodied intelligence industry to academic frontiers. The close coupling with conference names indicates that enterprises actively require practitioners to stay current with the latest research.

These three topics constitute a complete R&D cycle from algorithmic innovation to engineering validation to frontier monitoring, reflecting the systematic talent demands of enterprises transitioning from technological exploration to industrial implementation.

4.1.3. Simulation, Control and Deployment

The three topics of responsibility are: Simulation Loop Construction (159 entries), Motion Control (347 entries) and Model Deployment and Inference Optimization (134 entries); the total is 640 entries. The three subjects involve the whole chain of virtual simulation training to actual operation and edge deployment.

Simulation Loop Construction involves building high-fidelity simulation environments using platforms such as Isaac Sim, MuJoCo, and Gazebo, as well as designing Sim2Real transfer strategies. Keywords such as "high fidelity", "Sim2Real" and "rendering" suggest that one of the main aims is to reduce the difference between simulation and reality. A simulation environment can reduce the cost and risk of debugging physical robots, but practitioners should still have both the skill of simulation modeling and real-robot debugging.

Motion Control involves grasping operations, motion planning, obstacle avoidance and dynamics control. With 347 entries, it is the second-largest topic in this category and links the simulation-trained policies to the actual operation of robots. The keywords "reinforcement learning", "kinematics" and "dynamics" show that the control of current motion is moving towards learning-based methods and away from traditional model-based ones.

Model Deployment and Inference Optimization are concerned with the engineering problems of putting trained models on resource-constrained edge devices. TensorRT acceleration, GPU resource scheduling, model quantization and adaptation to embedded platforms are commonly required. Although the number (134) is relatively small, it is necessary for the real-time operation of intelligent systems on edge devices.

Together, these three topics form a complete engineering pipeline from virtual simulation to physical execution to edge deployment, reflecting the industry's need for full-stack capabilities in embodied intelligence system development.

4.1.4. System Integration and Testing

The only responsibility topic in this category is Testing, Integration and Commissioning (855 entries), which is the largest topic. This topic comprises the whole process of integrating, debugging and deploying all the software and hardware modules of embodied intelligent systems, such as setting up communication protocols under ROS2 (Robot Operating System 2), system integration in Linux environments, and joint-level debugging and testing of modules. Among all responsibility topics, this one has the

highest number of entries; therefore, in the field of embodied intelligent manufacturing, competence in a single module is no longer sufficient to meet the demands of the industry, and now there is a high demand for compound talents with system-level integration abilities. The keywords "ROS2", "Linux", "debugging" and "testing" show that system integration work relies heavily on the ecosystem of the robot operating system and Linux; therefore, good debugging and testing skills are also necessary. The fact that this subject is larger than all others indicates that the industry urgently needs professionals who can integrate heterogeneous subsystems and build a working system.

4.2. Responsibility-Skill Coupling Analysis

Association rule mining was employed to find associations among the responsibility topics and skill keywords. The four groups of responsibilities and their contents are as follows:

4.2.1. Perception and Data Engineering

The coupling patterns of the three topics in the Perception and Data Engineering category are shown in Tables 3-5.

As shown in Table 3, the coupling hierarchy of 3D Perception Modeling (25 rules) is steep. Among the single-skill rules, Detection (lift 7.037) and Segmentation (lift 6.954) are dominant, more than double that of generic Vision (lift 3.068); therefore, 3D perception can be regarded as a problem of detection and segmentation. The two-skill combination of Detection + Segmentation (lift 9.015) is the highest among all 128 rules and shows that this topic has good synergy. SLAM (lift 4.091) and 3D (lift 4.012) are a secondary level of spatial reasoning. The combination of the three skills: Algorithm + Framework + Vision (lift 6.818) and C++ + Framework + Vision (lift 6.027) are two complementary ways to reach the same foundation of framework and vision. The increase in lift with the addition of skills is monotonic: Vision (3.068) < Framework + Vision (5.071) < Algorithm + Framework + Vision (6.818), and the mean lift of multi-skill (5.36) is slightly higher than that of single-skill (5.03).

Table 3. Association Rules for 3D Perception Modeling

Keywords	Support	Confidence	Lift
Detection	0.024	0.279	7.037
Segmentation	0.021	0.246	6.954
SLAM	0.026	0.295	4.091
3D	0.026	0.295	4.012
Vision	0.038	0.443	3.068
Detection, Segmentation	0.02	0.23	9.015
Detection, Python	0.018	0.213	7.175
C++, SLAM	0.018	0.213	5.795
Python, SLAM	0.018	0.213	5.58
Framework, Vision	0.03	0.344	5.071

Vision, Computer Vision	0.018	0.213	5.022
Automation, SLAM	0.02	0.23	4.772
Algorithm, SLAM	0.023	0.262	4.636
Algorithm, Vision	0.032	0.377	4.518
Deep Learning, Vision	0.021	0.246	4.346
ROS, Vision	0.018	0.213	4.185
PyTorch, Vision	0.026	0.295	3.936
C++, Vision	0.023	0.262	3.863
Algorithm, Framework, Vision	0.028	0.328	6.818
C++, Framework, Vision	0.018	0.213	6.027
Algorithm, PyTorch, Vision	0.021	0.246	5.795
Algorithm, C++, Vision	0.021	0.246	5.433
Python, Framework, Vision	0.024	0.279	5.325
Algorithm, Python, Vision	0.027	0.312	5.243
Automation, Framework, Vision	0.018	0.213	4.708

Table 4 shows the 18 rules of Multi-Sensor Integration and Fusion. IMU (lift 5.041) and Camera (lift 4.332) are better predictors than the generic Sensor (lift 3.599); although the generic Sensor appears in 16 out of 18 rules and has the highest support (0.051), it has the lowest lift and thus low discriminatory ability. The two skill combinations with the highest lifts are Sensor + IMU and Sensor + Linux, at 5.198 and 5.050, respectively. Sensor + C++ and Sensor + Python also have relatively high lifts (4.386 and 4.229), indicating that both languages are essential for sensor driver development. Combining the three skills of C++ + Sensor + Linux (lift 5.378) and Automation + Sensor + C++ (lift 4.512) can cover all parts of the sensor integration stack.

Table 4. Association Rules for Multi-Sensor Integration and Fusion

Keywords	Support	Confidence	Lift
IMU	0.023	0.235	5.041
Camera	0.021	0.221	4.332
Sensor	0.051	0.529	3.599
Sensor, IMU	0.021	0.221	5.198
Sensor, Linux	0.024	0.25	5.05
Debugging, Sensor	0.024	0.25	4.419
C++, Sensor	0.038	0.397	4.386
Python, Sensor	0.034	0.353	4.229
Hardware, Sensor	0.021	0.221	3.999
ROS2, Sensor	0.026	0.265	3.982
Automation, Sensor	0.041	0.426	3.916
C++, Sensor, Linux	0.021	0.221	5.378
Automation, Sensor, C++	0.032	0.338	4.512
Python, Automation, Sensor	0.028	0.294	4.424
Sensor, Automation, ROS	0.026	0.265	4.352
Sensor, C++, ROS	0.028	0.294	4.332

Python, C++, Sensor	0.031	0.324	4.316
ROS2, Python, Sensor	0.024	0.25	4.311

Table 5 shows the rules for Multimodal Data Governance, and the sparsest topic has only 2 rules: Data Processing (lift 2.320) and PyTorch + Python (lift 1.330). The latter is just slightly above the independence baseline of 1.0, so it does not have much discriminative ability for the role of data governance. There are no particular rules, so the responsibility for data governance falls on the general-purpose data processing skills and not on specialised toolchains; it can be incorporated into the data engineering work at all levels of the system.

Table 5. Association Rules for Multimodal Data Governance

Skill Keywords	Support	Confidence	Lift
Data Processing	0.061	0.207	2.32
PyTorch, Python	0.064	0.216	1.33

4.2.2. Algorithm and Model Development

The coupling patterns of the three topics in the Algorithm and Model Development category are shown in Tables 6-8.

As shown in Table 6, the coupling structure of Multimodal Embodied Modeling (18 rules) is highly related to academic research. Diffusion (lift 3.011) is the top of the single-skill hierarchy, and it is followed by CVPR (2.979), NeurIPS (2.948), VLM (2.709), Transformer (2.676), and ICRA (2.578). Recently, diffusion models have begun to be used in the problems of embodied artificial intelligence for generative models. The three-skill rule PyTorch + Framework + Multimodal (lift 2.818) refers to the technical stack.

Table 6. Association Rules for Multimodal Embodied Modeling

Keywords	Support	Confidence	Lift
Diffusion	0.061	0.23	3.011
CVPR	0.074	0.278	2.979
NeurIPS	0.065	0.246	2.948
VLM	0.068	0.257	2.709
Transformer	0.065	0.246	2.676
ICRA	0.064	0.241	2.578
NeurIPS, CVPR	0.057	0.214	3.517
Paper, Multimodal	0.058	0.219	3.1
Paper, CVPR	0.064	0.241	2.933
Paper, NeurIPS	0.055	0.209	2.891
Deep Learning, Multimodal	0.062	0.235	2.772
PyTorch, Multimodal	0.079	0.3	2.714
Large Model, Framework	0.062	0.235	2.683
Large Model, Model	0.071	0.267	2.662
Large Model, PyTorch	0.062	0.235	2.64

Large Model, Deep Learning	0.055	0.209	2.633
Paper, PyTorch	0.068	0.257	2.63
PyTorch, Framework, Multimodal	0.058	0.219	2.818

Table 7 shows the rules for Data Loop and Evaluation (7 rules), and all lifts are below 2.0; the single-skill rules are closely grouped between 1.68 and 1.87 (spread of 0.194). The single-skill hierarchy is: Training (lift 1.874), Large-scale (1.778), MuJoCo (1.741) and Multimodal (1.680); none of them stand out significantly among the others. The two-skill rules, Intelligence + Model (lift 1.980) and Model + Multimodal (1.909), show that the scope of evaluation needs to be extended to include model understanding in addition to black-box tests. Coupling with MuJoCo (1.741) and Model + Automation (1.822) shows that the main evaluation platform is physics simulation.

Table 7. Association Rules for Data Loop and Evaluation

Skill Keywords	Support	Confidence	Lift
Training	0.076	0.278	1.874
Large-scale	0.057	0.206	1.778
MuJoCo	0.061	0.222	1.741
Multimodal	0.1	0.366	1.68
Intelligence, Model	0.062	0.227	1.98
Model, Multimodal	0.062	0.227	1.909
Model, Automation	0.062	0.227	1.822

Table 8 shows the rules for Frontier Technology Tracking (12 rules). Ph.D. (lift 2.605) is the strongest single-skill predictor, surpassing Paper (2.161) by 20.6%—the only topic where a formal qualification, rather than a technical skill, is the top predictor. The two-skill rules Paper + Framework (lift 2.609) and Paper + PyTorch (2.546) indicate that technology tracking combines publication literacy with hands-on experimentation. PyTorch + Reinforcement Learning (2.425) and Intelligence + PyTorch (2.374) reflect the actively tracked research domains. The three-skill rule Intelligence + PyTorch + Framework (2.448) encapsulates the complete technology tracking stack. Multi-skill mean lift (2.40) exceeds single skill (2.14) by 12.1%, the largest relative increase in this category.

Table 8. Association Rules for Frontier Technology Tracking

Keywords	Support	Confidence	Lift
Ph.D.	0.05	0.217	2.605
Paper	0.088	0.385	2.161
Deep Learning	0.082	0.36	1.901
PyTorch	0.095	0.416	1.886
Paper, Framework	0.058	0.255	2.609
Paper, PyTorch	0.057	0.248	2.546

PyTorch, Reinforcement Learning	0.052	0.23	2.425
Intelligence, PyTorch	0.057	0.248	2.374
Paper, Python	0.055	0.242	2.314
Deep Learning, Framework	0.069	0.304	2.265
Intelligence, Reinforcement Learning	0.051	0.224	2.196
Intelligence, PyTorch, Framework	0.048	0.211	2.448

These three topics represent distinct facets of the algorithm development pipeline: Multimodal Embodied Modeling is driven by top-tier conferences and state-of-the-art model architectures, Data Loop and Evaluation centers on the engineering practices of model testing and automation, and Frontier Technology Tracking is distinguished by its emphasis on research publications and formal academic credentials.

4.2.3. Simulation, Control and Deployment

The coupling patterns for the three topics within this category are reported in Tables 9–11.

Table 9 presents the rules for Simulation Loop Construction (15 rules). Physics (lift 3.571), Simulation Platform (3.422), and Dynamics (3.209) all exceed the generic Simulation (2.940), with Physics surpassing it by 21.5%. Simulation, despite having the highest support (0.075), has the lowest lift, which mirrors the Sensor pattern. The two-skill combinations Sim + Simulation (lift 3.535), Isaac + Simulation (3.399), and Isaac + C++ (3.390) couple platform expertise with systems programming. Both the three-skill rule Sim + Python + Simulation (3.764) and Isaac + Sim + Simulation (3.491) cover all the simulation skill requirements. C++ is consistently used at all levels (Isaac + C++: 3.390; C++ + Simulation: 3.317; Sim + C++: 3.273).

Table 9. Association Rules for Simulation Loop Construction

Keywords	Support	Confidence	Lift
Physics	0.034	0.222	3.571
Simulation Platform	0.032	0.213	3.422
Dynamics	0.035	0.232	3.209
Simulation	0.075	0.491	2.94
Sim, Simulation	0.038	0.25	3.535
Isaac, Simulation	0.038	0.25	3.399
Isaac, C++	0.041	0.268	3.39
Isaac, Python	0.05	0.324	3.321
C++, Simulation	0.054	0.352	3.317
Sim, C++	0.042	0.278	3.273
Sim, Python	0.048	0.315	3.226
Sim, Python, Simulation	0.032	0.213	3.764
Isaac, Sim, Simulation	0.034	0.222	3.491
Isaac, Python, C++	0.041	0.268	3.452

Python, Automation, Simulation	0.034	0.222	3.343
--------------------------------	-------	-------	-------

As shown in Table 10, Motion Control (11 rules) is centered on the ROS ecosystem and reinforcement learning. The single-skill hierarchy—Control (lift 1.890), Reinforcement Learning (1.788), ROS (1.741), and Vision (1.644)—reflects the coexistence of traditional control and learning-based approaches. Reinforcement Learning + ROS (lift 2.277) achieves the highest lift, a 30.8% increase over single-skill ROS, indicating the synergy of learning-based control and robot operating systems. Algorithm + ROS (2.012) and Automation + Reinforcement Learning (1.921) further support this pattern. The three-skill rules Algorithm + Python + ROS (2.161) and Algorithm + Python + Framework (2.017) encapsulate the motion control stack. Python consistently couples across multi-skill combinations (Python + ROS: 1.908; Algorithm + Python + ROS: 2.161).

Table 10. Association Rules for Motion Control

Skill Keywords	Support	Confidence	Lift
Control	0.103	0.34	1.89
Reinforcement Learning	0.115	0.377	1.788
ROS	0.127	0.419	1.741
Vision	0.072	0.237	1.644
Reinforcement Learning, ROS	0.064	0.209	2.277
Algorithm, ROS	0.074	0.242	2.012
Automation, Reinforcement Learning	0.074	0.242	1.921
Python, ROS	0.108	0.354	1.908
Reinforcement Learning, Framework	0.065	0.214	1.891
Algorithm, Python, ROS	0.065	0.214	2.161
Algorithm, Python, Framework	0.065	0.214	2.017

Table 11 shows the 15 rules for Model Deployment and Inference Optimization. Transformer (2.670) is the highest in the single-skill ranking, and below it are TensorFlow (2.310), PyTorch (2.142) and Deployment (2.081). Deployment + Model (lift 2.868) is the highest, and it is 37.8% higher than single-skill Deployment. Optimization is carried out in C++ at all levels (PyTorch + C++: 2.511; C++ + Multimodal: 2.594). The three-skill rule is PyTorch + Model + Framework (2.646) and Python + C++ + Multimodal (2.594); these are the deployment stacks. Multimodal consistently co-occurs with deployment-related skills across multiple rule combinations (e.g., C++ + Multimodal: lift 2.594; PyTorch + Multimodal: lift 2.554), indicating that multimodal models are a common deployment target in this domain.

Table 11. Association Rules for Model Deployment and Inference Optimization

Keywords	Support	Confidence	Lift
Transformer	0.038	0.246	2.67
TensorFlow	0.032	0.209	2.31
PyTorch	0.074	0.473	2.142
Deployment	0.064	0.409	2.081
Deployment, Model	0.041	0.264	2.868
Deployment, PyTorch	0.032	0.209	2.64
C++, Multimodal	0.032	0.209	2.594
PyTorch, Multimodal	0.044	0.282	2.554
PyTorch, C++	0.035	0.227	2.511
Multimodal, Framework	0.04	0.254	2.5
Model, Multimodal	0.045	0.291	2.448
Algorithm, Multimodal	0.034	0.218	2.373
Model, Deep Learning	0.035	0.227	2.363
PyTorch, Model, Framework	0.04	0.254	2.646
Python, C++, Multimodal	0.032	0.209	2.594

Together, the three topics span the full chain from virtual simulation to physical execution to edge deployment, each demanding foundational domain knowledge: Physics and Dynamics for simulation, Control theory for motion control, and Model architecture understanding for deployment.

4.2.4. System Integration and Testing

Table 12 shows the rules for Testing, Integration and Commissioning (5 rules), with all lifts below 1.7. Linux (lift 1.657) is a stronger predictor than ROS2 (1.420), underscoring the foundational role of systems-level proficiency. The two-skill rules ROS2 + C++ (1.486) and ROS2 + Python (1.486) exhibit identical lifts, indicating language-agnostic coupling. The three-skill rule Python + Automation + C++ (1.278) captures the multi-language, automation-oriented nature of integration testing. Notably, single-skill mean lift (1.54) exceeds multi-skill (1.42)—the only topic where skill combination reduces discriminative power, suggesting that integration competence is characterized by breadth of exposure rather than depth in any single toolchain.

Table 12. Association Rules for Testing, Integration and Commissioning

Keywords	Support	Confidence	Lift
Linux	0.134	0.262	1.657
ROS2	0.147	0.287	1.42
ROS2, C++	0.122	0.238	1.486
ROS2, Python	0.122	0.238	1.486
Python, Automation, C++	0.123	0.24	1.278

4.3. Cross-Category Coupling Patterns

Beyond the intra-category analysis, several cross-category coupling patterns emerge from the association rule results.

First, Python and PyTorch appear as high-lift skills across nearly all categories, forming the common skill substrate of the embodied intelligence technology stack. Second, C++ emerges as a critical skill in hardware-adjacent categories (Sensor, lift 4.386; Simulation, lift 3.317), reflecting the performance requirements of real-time systems. Third, the ROS/ROS2 ecosystem spans multiple categories, from sensor integration to motion control to system integration, confirming its role as the central middleware for embodied intelligent systems. Fourth, the coupling between conference venues (CVPR, NeurIPS, ICRA) and both Multimodal Embodied Modeling and Frontier Technology Tracking indicate a strong and bidirectional industry-academia connection. Finally, the sparse rule sets for Multimodal Data Governance and System Integration suggest that these responsibilities rely on general-purpose engineering skills rather than specific toolchains. Collectively, these patterns reveal that embodied intelligent manufacturing demands a T-shaped skill profile: broad proficiency in Python and PyTorch across all categories, combined with deep specialization in category-specific tools, alongside continuous knowledge updating driven by close academic ties.

5. Conclusion

5.1. Discussion

The results of this study reveal a systematic structure of responsibility topics and skill demands in embodied intelligent manufacturing. Through the improved BERTopic method, 10 responsibility topics were identified and categorized into four clusters: Perception and Data Engineering (Multimodal Data Governance, Multi-Sensor Integration and Fusion, 3D Perception Modeling), Algorithm and Model Development (Multimodal Embodied Modeling, Data Loop and Evaluation, Frontier Technology Tracking), Simulation, Control and Deployment (Simulation Loop Construction, Motion Control, Model Deployment and Inference Optimization), and System Integration and Testing (Testing, Integration and Commissioning). Association rule mining produced 128 filtered rules linking these topics to skill keywords, with lift values ranging from 1.278 to 9.015. The cross-category analysis further revealed that Python and PyTorch form the common skill substrate across all categories, C++ is critical in hardware-adjacent domains, and the ROS/ROS2 ecosystem serves as the central middleware spanning Perception, Control, and System Integration categories. The coupling between conference venues (CVPR, NeurIPS, ICRA) and both Multimodal Embodied Modeling and Frontier Technology Tracking indicate a strong and bidirectional industry-academia connection. Collectively, these patterns indicate that embodied intelligent manufacturing requires a T-shaped skill profile: horizontal breadth in Python and PyTorch, coupled with vertical depth in domain-specific tools.

The three main contributions of this paper are:

- A scalable analytical framework for job advertisement – text. This study constructs a complete analytical pipeline covering data acquisition, preprocessing, topic modeling, and association rule mining. In the data acquisition stage, Playwright was used to simulate browser behavior for crawling job advertisements, and the Qwen3.7-max large language model was employed to extract job responsibilities and skill requirements. At the stage of topic modeling, an improved BERTopic method that integrates GSDMM and semantic representation has been introduced to improve the accuracy of short-text topic modeling by combining multi-topic GSDMM similarity calculation with semantic vector similarity. In the association rule mining step, the Apriori algorithm was used to find association rules among responsibility topics and skill keywords, and the elbow method was employed to determine the threshold. The above framework is relatively simple and can be used for analysing job advertisement text in other areas.
- Based on the improved BERTopic model, this paper identifies ten responsibility topics in embodied intelligent manufacturing positions, which are further organized into four functional groups. Association rule mining reveals significant couplings between responsibility topics and skill requirements. Each cluster exhibits distinct skill signatures: the Perception and Data Engineering group is dominated by domain-specific algorithmic skills, Algorithm and Model Development is strongly coupled with academic research indicators, Simulation, Control and Deployment centers on simulation platforms and learning-based control methods, and System Integration and Testing is anchored by robotics ecosystem tools. These identified linkages can serve as a reference index for talent demand in intelligent manufacturing.
- Support for Manufacturing Enterprises in Strengthening Talent Allocation Adaptability and Technology Deployment Resilience. The responsibility-skill coupling relationships in this paper can help manufacturing enterprises understand the skill demands of all sectors in the industry and optimise talent allocation and training. For instance, in the case of Perception and Data Engineering, the company needs to find people who are knowledgeable about sensor calibration, detection and segmentation algorithms; for Algorithm and Model Development, one should hire individuals skilled in multimodal architectures who can keep up with the latest developments in academia; for Simulation, Control and Deployment, practitioners must be able to conduct model simulations and debug real robots; finally, a strong foundation in Linux and ROS2 at the system level is required for System Integration and

Testing. At the same time, these results can help colleges and universities adjust their courses and train more interdisciplinary students according to the requirements of the industry.

5.2. Limitations and Future Directions

This paper has the following shortcomings:

- **Limited Data Volume.** A total of 749 job advertisements were collected at one time. As embodied intelligent manufacturing is a relatively new field, the demands on talent at present will change in the future. In the future, we will continuously monitor the recruitment data from various periods to understand how the demands for skills in this rapidly changing industry have evolved.
- **High computational cost and limited scalability of GSDMM similarity computation.** The GSDMM similarity computation involves two computationally costly stages. First, a grid search over GSDMM hyperparameters (α , β) must be performed to identify the optimal configuration, and then multiple GSDMM training runs under different topic numbers to generate the topic probability distributions for similarity calculation. This multi-stage pipeline incurs substantial time cost. Second, the final similarity matrix needs to calculate all-pairs similarity among sentences, resulting in $O(N^2)$ memory complexity, where N is the number of responsibility entries. For larger datasets, this quadratic growth makes the memory footprint prohibitive. Furthermore, the pipeline lacks incremental learning capability: any new data requires full retraining of all stages from scratch. Future work could explore more efficient similarity computation methods to reduce the computation cost and enable the framework to scale to larger datasets.
- **This study does not address potential bias and fairness issues in the analyzed job advertisement texts, which prior work has shown may affect hiring outcomes.** Future work could incorporate bias detection components into the proposed framework.
- **Qwen3.7-max was used for information extraction and Qwen3-Embedding for semantic vector generation in this paper.** A two-stage verification has been added to reduce the bias of a single model, but it is not practical or economical to use commercial APIs. In the future, adopting open-source LLMs could serve as a viable alternative."
- **The dataset in this paper does not include scenario labels (e.g., assembly, inspection, flexible production), preventing a division of responsibility-skill coupling patterns by manufacturing scenario in the analysis.** The mapping of responsibility clusters to scenarios in Section 4.1 is therefore grounded in domain knowledge rather than empirical data. To improve the accuracy of the responsibility-scenario map in future

work, additional data with explicit scenario labels should be gathered, and enterprise surveys that directly link job responsibilities to production scenarios could be carried out.

In addition, the following directions for future research can be considered: (1) Applying the proposed analytical framework to other fields to check for generalizability. (2) Introducing dynamic analysis methods to track fluctuations in skill requirements and anticipate shifts in future skill demand; and (3) Incorporating actual talent allocation data of enterprises to build talent demand-supply matching models that can provide more direct support for corporate talent strategy decisions.

References

- [1] Zhu, J., S. Huang, P. Wang, et al., *A review on large language models for industrial embodied intelligence*. Advanced Engineering Informatics, 2026. 73: p. 104602.
- [2] Fang, L., J. Shi, L. Wu, et al., *Perspectives and Prospects on Embodied Intelligence-Empowered Smart Manufacturing*. Journal of Intelligent Manufacturing, 2026: p. 1-20.
- [3] Xu, J., Q. Sun, Q.-L. Han, et al., *When Embodied AI Meets Industry 5.0: Human-Centered Smart Manufacturing*. Ieee-caa Journal of Automatica Sinica, 2025. 12(3): p. 485-501.
- [4] Al-Asfour, A. and Y. Zhao, *Bridging the skills gap divide in manufacturing: perspectives from industry leaders*. Industrial and Commercial Training, 2024. 56(1): p. 78-90.
- [5] Menezes, P.C., K.M. Sakiyama, R.M. Marcacini, et al., *SKORE: Skill-oriented extraction engine from job advertisements using large language models*. MethodsX, 2026. 16: p. 103907.
- [6] Gurcan, F., A. Soylu, and A.Q. Khan, *Towards a Sustainable Workforce in Big Data Analytics: Skill Requirements Analysis from Online Job Postings Using Neural Topic Modeling*. Sustainability, 2025. 17(20).
- [7] Kang, P.S., R. Enstroem, B. Bhawna, et al., *A text mining study of competencies in modern supply chain management with skillset mapping*. Supply Chain Analytics, 2025. 10: p. 100117.
- [8] Gavrilescu, M., F. Leon, and A.-A. Minea, *Techniques for Transversal Skill Classification and Relevant Keyword Extraction from Job Advertisements*. Information, 2025. 16(3): p. 167.
- [9] Grootendorst, M., *BERTopic: Neural topic modeling with a class-based TF-IDF procedure*. CoRR, 2022. abs/2203.05794.
- [10] Yin, J. and J. Wang, *A Dirichlet Multinomial Mixture Model-based Approach for Short Text Clustering*, in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2014, ACM: New York, NY, USA. p. 233-242.
- [11] Xia, D., P. Xu, G. Han, et al., *Emerging Perspectives on Embodied Intelligence in Future Smart Manufacturing*. Journal of Industrial Information Integration, 2026. 50: p. 101020.
- [12] Li, S., C. Liu, J. Zhao, et al., *Physics-informed embodied intelligence in the foundation model era: Advancing robot manipulation for smart manufacturing*. Advanced Engineering Informatics, 2026. 71: p. 104370.

- [13] Ding, K., Q. Mao, Y. Zhang, et al., *Review and perspectives on multimodal perception, mutual cognition, and embodied execution for human–robot collaboration in Industry 5.0*. Robotics and Computer-Integrated Manufacturing, 2026. 101: p. 103280.
- [14] Pejic-Bach, M., T. Bertoncel, M. Meško, et al., *Text mining of industry 4.0 job advertisements*. International Journal of Information Management, 2020. 50: p. 416-431.
- [15] Schedlbauer, J., G. Raptis, and B. Ludwig, *Medical informatics labor market analysis using web crawling, web scraping, and text mining*. Int J Med Inform, 2021. 150: p. 104453.
- [16] Verma, A., K. Lamsal, and P. Verma, *An investigation of skill requirements in artificial intelligence and machine learning job advertisements*. Industry and Higher Education, 2022. 36(1): p. 63-73.
- [17] Wang, R., Q. Chen, Y. Wang, et al., *JobViz: Skill-driven visual exploration of job advertisements*. Visual Informatics, 2024. 8(3): p. 18-28.
- [18] Chaturvedi, S., K. Mahajan, and Z. Siddique, *Gendered language in job ads and applicant behavior: Evidence from India*. Labour Economics, 2025. 96: p. 102726.
- [19] Hu, S., J.A. Al-Ani, K.D. Hughes, et al., *Balancing Gender Bias in Job Advertisements with Text-Level Bias Mitigation*. FRONTIERS in BIG DATA, 2022. 5.
- [20] Zhong, A., R. Qiu, X. Zhang, et al., *Characterizing the Job-task-skill Pattern of Job Requirements with Job Advertisement Mining*, in *Proceedings of the 2023 6th International Conference on Computers in Management and Business*. 2023. p. 139-147.
- [21] Zhu, Y. and Y. Liu, *Gibbs-BERTopic: A Hybrid Approach for Short Text Topic Modeling*. IEEE Access, 2025. 13: p. 49162–49173.
- [22] Bylina, B. and A. Antończak, *Analysis of end-to-end test automation tools based on the examples of Selenium WebDriver and Playwright*, in *Position Papers of the 19th Conference on Computer Science and Intelligence Systems*. 2024. p. 1–8.
- [23] QwenTeam. *Qwen3.7*. 2026; Available from: <https://qwen.ai/blog?id=qwen3.7>.
- [24] DeepSeek-AI. *DeepSeek-V3*. 2024; Available from: <https://api-docs.deepseek.com/news/news1226/>.
- [25] QwenLM. *Qwen3-Embedding*. 2025; Available from: <https://github.com/QwenLM/Qwen3-Embedding>.
- [26] Agrawal, R. and R. Srikant. *Fast Algorithms for Mining Association Rules in Large Databases*. in *Proceedings of the 20th International Conference on Very Large Data Bases*. 1994. Morgan Kaufmann Publishers Inc.
- [27] Satopaa, V., J.R. Albrecht, D.E. Irwin, et al., *Finding a "kneedle" in a Haystack: Detecting Knee Points in System Behavior*. ICDCS Workshops, 2011: p. 166-171.