

Graph Neural Network-Based Cascading Disruption Path Identification in Multi-Tier Rare Earth Processing Networks

Yanhuan Chen¹, Jiacheng Hu²

¹ Master of Engineering, Dartmouth College, NH, USA

² Master's Degree in Information Technology, University of New South Wales, Australia

Abstract

The United States confronts acute strategic vulnerability across rare earth supply chains, with import dependency on China exceeding 80% at most processing stages. This concentration risk is amplified by the sequential multi-tier structure of rare earth processing, in which primary mining, oxide separation, metal refining, alloy manufacturing, and permanent magnet production form a strictly ordered transformation hierarchy. A disruption at any single tier propagates to adjacent stages in nonlinear and often asymmetric ways, yet existing quantitative frameworks remain insufficient for tracing these cascading pathways at the tier level. This paper develops a graph neural network (GNN) framework applied to a heterogeneous multi-tier network representing rare earth processing flows across major producing and consuming nations. The framework integrates bilateral trade data from the BACI International Trade Database, processing capacity statistics from the U.S. Geological Survey (USGS) and the International Energy Agency (IEA), and the Caldara-Iacoviello Geopolitical Risk Index as node-level features. A relational graph convolutional network (R-GCN) simulates the propagation of disruption signals across processing tiers, and a graph attention module ranks cascading pathways by cumulative vulnerability scores. Experiments conducted on 34 documented disruption events from 2000 to 2024 identify oxide separation as the dominant structural bottleneck tier, and ablation analysis confirms that heterogeneous edge treatment and geopolitical risk features each contribute independently to model performance. The framework quantifies how alternative sourcing corridors alter network-level vulnerability across four geopolitical disruption scenarios.

Keywords: rare earth supply chain; graph neural networks; cascading disruption; supply chain resilience

1. Introduction

1.1 Background and Motivation

The global rare earth supply chain exhibits a degree of geographic concentration that distinguishes it from virtually all other critical industrial material systems. China's dominant position extends well beyond primary mining: it controls over 90% of global rare earth oxide separation capacity and accounts for approximately 94% of sintered permanent magnet production, a configuration that gives a single nation decisive leverage over every downstream industry dependent on these materials [1]. The United States does not simply depend on a foreign supplier for a raw commodity—it depends on that supplier for nearly every transformation stage between ore and functional end component. The implications are substantial for semiconductor manufacturing, precision defense systems, wind turbine generators, and electric vehicle drivetrains, all of which require permanent magnets containing neodymium and dysprosium whose processing chain runs almost entirely through Chinese facilities.

This vulnerability became visible in 2010, when China imposed export quotas on rare earth oxides, triggering a supply shock that raised international spot prices by more than 700% within eighteen months and forced downstream manufacturers in Japan, Europe, and North America to scramble for alternatives that did not exist at scale. The episode exposed the inherent fragility of a processing architecture offering no meaningful substitution at any intermediate stage. Subsequent years brought repeated demonstrations of the same structural problem. China's April 2025 escalation—adding seven rare earth elements and rare earth permanent magnets to its strategic export control list following earlier restrictions on gallium, germanium, and graphite imposed in 2023—represents a deliberate and progressive application of supply chain leverage as a geopolitical instrument [2]. The USGS reported that U.S. net import reliance for rare earth compounds and metals reached 100% in 2024, with China accounting for 70% of processed rare earth imports and controlling essentially all globally deployed oxide separation capacity [3].

The challenge is not solely one of sourcing. The rare earth processing chain is a sequential system in which the output of each transformation stage is the necessary input for the next. A disruption that constrains capacity at one tier does not simply reduce throughput at that stage; it propagates forward and backward through the network in ways that cannot be anticipated from aggregate supply-demand analysis alone. Despite the growing

policy attention to critical mineral supply chain resilience, the connection between measurable geopolitical risk signals and specific disruption pathways within multi-tier mineral processing networks remains analytically underspecified in the quantitative literature [4].

1.2 Research Objectives and Contributions

A. Research Questions

This paper addresses two interrelated research questions. The first concerns structural propagation mechanics: in a multi-tier rare earth processing network, through which specific pathways does an initial disruption at one node or tier transmit to other processing stages, and what determines the magnitude and direction of that transmission? The second concerns vulnerability identification: which particular processing tiers and country-tier combinations constitute the most critical bottlenecks in the network, as measured by their capacity to amplify disruption signals under different geopolitical scenarios?

B. Key Contributions

This paper makes three contributions to the literature on critical mineral supply chain analysis. The first is the construction of a heterogeneous multi-tier directed graph representing the rare earth processing network across 47 country-tier node pairs and five sequential processing stages, integrating publicly available trade flow and processing capacity data from institutional sources. The second is the application of an R-GCN to simulate disruption propagation through this heterogeneous structure using a message-passing mechanism adapted to capture tier-level transmission asymmetries across two distinct edge relation types. The third is the introduction of a graph attention-based path ranking procedure that produces interpretable vulnerability scores for cascading pathways, enabling structured comparative analysis across geopolitical disruption scenarios. The scope is deliberately bounded to rare earth processing in order to preserve empirical tractability and maintain internal validity within a single mineral system.

2. Related Work

2.1 Critical Mineral Supply Chain Risk Assessment

A. Traditional Network Analysis Approaches

Quantitative assessment of critical mineral supply chain resilience has drawn extensively on complex network theory, with research groups applying node centrality measures, clustering coefficients, and topological indicators to bilateral trade networks. Sprecher et al. constructed a resilience framework for material supply chains using the 2010 rare earth crisis as a validation case, demonstrating that network topology metrics could differentiate between resilient and fragile supply configurations across multiple disruption scenarios, and establishing that response, recovery, and rapidity are the three operationalizable dimensions of material supply chain resilience [5]. A subsequent study by the same group extended this approach to develop novel resilience indicators specifically calibrated to critical material supply chains, confirming that the multi-tier processing structure of rare earths introduces disruption amplification dynamics that simpler aggregate graph metrics cannot adequately represent, and establishing tier-explicit network analysis as a methodological necessity rather than a refinement [6]. These contributions remain foundational to the field, yet they rely on static network snapshots and implicitly linear propagation assumptions that constrain their ability to represent the nonlinear dynamics observable in empirical disruption events.

B. AI and Machine Learning Approaches

The application of artificial intelligence to supply chain risk management has expanded considerably in the past decade. Baryannis et al. conducted a systematic review cataloguing applications of machine learning, optimization, and simulation across supply chain risk identification, assessment, and mitigation tasks, identifying a persistent methodological gap in approaches capable of handling the interconnected and dynamic character of multi-tier supply networks where risk propagates along relational structures rather than within isolated firm-level processes [7]. Baskaran et al. subsequently provided empirical evidence that corporate investment in AI technologies is associated with a statistically significant reduction in exposure to critical mineral supply disruptions, with the largest effects concentrated among firms dependent on rare earth and battery mineral inputs—a finding that motivates the adoption of GNN-based methods specifically for this application domain [8].

2.2 Graph Neural Networks in Supply Chain Management

Graph neural networks have emerged as a structurally appropriate tool for supply chain analysis because supply chains are inherently relational systems in which risk propagates along explicitly defined inter-firm and inter-tier links. Kosasih and Brintrup demonstrated that a GNN-based link prediction approach could identify latent supply relationships not visible in procurement records, achieving substantially higher precision and recall than logistic regression and random forest baselines across a real automotive supply chain dataset,

and establishing the empirical viability of graph-based learning for supply chain structural inference [9]. The HKTGNN framework extended this direction by introducing hierarchical knowledge transfer across supply chain tiers, enabling the model to exploit structural regularities that repeat at different network levels and showing improved performance on supply chain risk classification tasks relative to flat GNN architectures that treat all nodes as structurally equivalent [10]. These works collectively establish the technical feasibility of GNN-based methods for multi-tier supply chain structures, though neither addresses the specific characteristics of mineral processing networks, including strictly unidirectional physical transformation flows and extreme geographic concentration at intermediate processing stages.

2.3 Geopolitical Risk and Mineral Supply Disruption

The quantification of geopolitical risk as a structured variable has been advanced by the Caldara and Iacoviello index, constructed from the frequency of geopolitical risk terms in major international newspapers and validated against investment decisions, bilateral trade flows, and financial market outcomes across multiple country contexts and time periods [11]. In the critical mineral domain, Liu et al. applied the DPSIR framework combined with machine learning techniques to assess the resilience of China's cobalt supply chain, incorporating geopolitical pressure as a driver-level variable and demonstrating that ML-augmented multidimensional assessment outperforms single-indicator approaches in capturing dynamic risk interactions that unfold across interconnected supply tiers [12]. A complementary bibliometric and text mining analysis by Liu et al. mapped the intellectual structure of the critical minerals supply chain security literature, identifying geopolitical risk quantification and multi-tier network modeling as two of the most underexplored intersections in the scholarly body, and calling for methods that bridge macroeconomic risk indicators with microstructural supply chain network analysis [13].

3. Methodology

3.1 Multi-Tier Processing Network Construction

The rare earth processing chain is modeled as a directed heterogeneous graph $G = (V, E, T_V, T_E)$, where V denotes the node set, E the directed edge set, T_V the set of node types corresponding to processing tiers, and T_E the set of edge types representing inter-tier trade flows and intra-tier competitive supply relationships. Five sequential processing tiers are defined: Tier 1 (Primary Mining), Tier 2 (Oxide Separation), Tier 3 (Metal Refining), Tier 4 (Alloy Manufacturing), and Tier 5 (Permanent Magnet Production). Each node v_i^k represents country i at processing tier k , included in the graph if the country-tier pair accounts for at least 1% of global throughput at that tier according to IEA and USGS capacity statistics [14]. The resulting network contains 47 country-tier nodes and 189 directed edges, with edge weights normalized by bilateral trade volumes derived from the UN Comtrade and BACI databases.

The network is structured to reflect the physical directionality of rare earth transformation: materials flow exclusively from lower-numbered tiers to higher-numbered tiers under normal operating conditions, and this directional constraint is encoded through separate edge type embeddings that prevent the GNN from conflating transformation-linked propagation with competitive reallocation dynamics. Intra-tier edges connect two nodes at the same tier if they export to a common set of downstream buyers, representing the competitive substitution relationships that determine whether an alternative supplier can absorb demand displaced by a disruption at a primary source node. Table 1 summarizes the network topology statistics across all five processing tiers, including node count, directed edge count, and the Herfindahl-Hirschman Index (HHI) computed from bilateral flow data as a structural concentration measure.

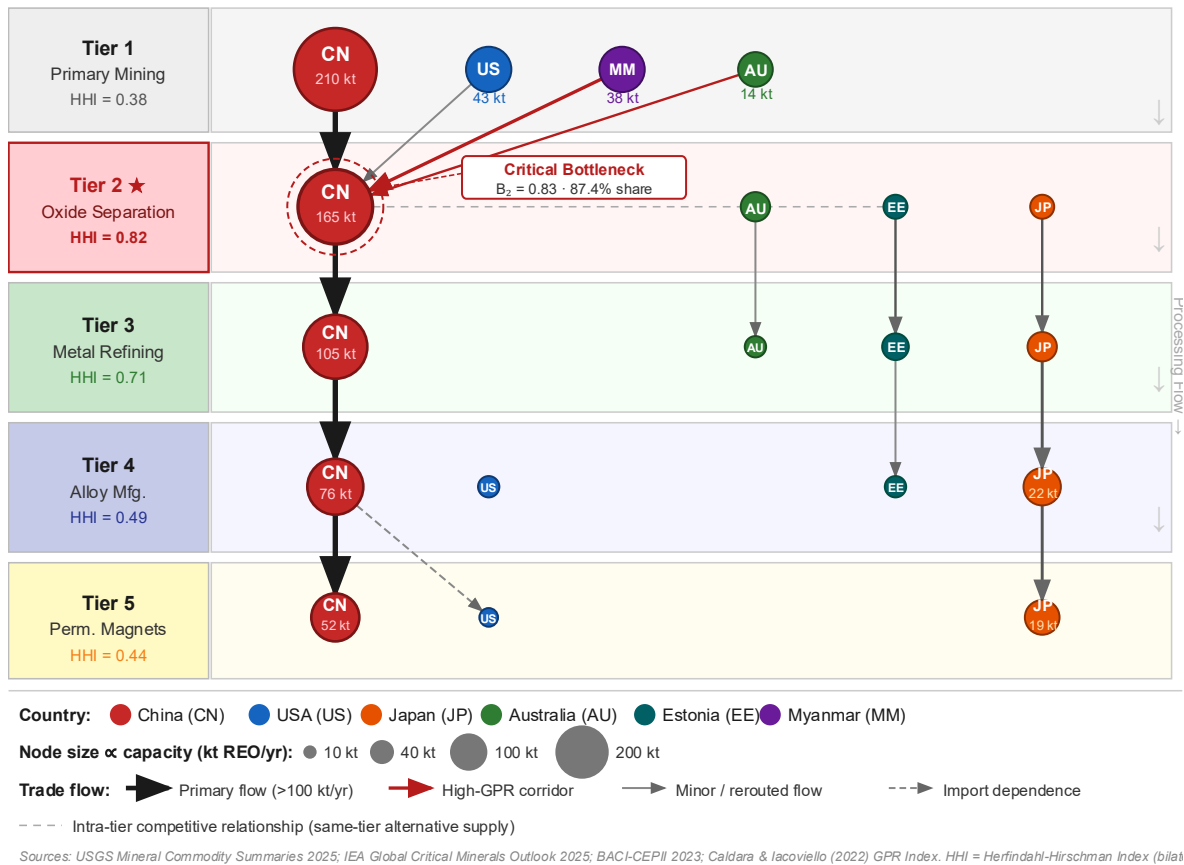
Table 1. Network topology statistics across five rare earth processing tiers (2023 baseline).

Tier	Tier Name	Node Count	Edge Count	HHI (Bilateral Flow)	Top Country Share (%)
1	Primary Mining	12	34	0.38	60.2 (China)
2	Oxide Separation	6	18	0.82	87.4 (China)
3	Metal Refining	8	24	0.71	79.3 (China)
4	Alloy Manufacturing	11	52	0.49	64.1 (China)

Note: HHI computed from BACI bilateral trade data (HS codes 2846, 7202.93, 8505.11). Top Country Share derived from USGS MCS 2025 [3] and IEA Global Critical Minerals Outlook 2025 [14]. HHI values approaching 1.0 indicate near-monopolistic concentration.

Node feature vectors x_i^k are assembled from three data streams: normalized processing capacity from the IEA Global Critical Minerals Outlook 2025 and USGS Mineral Commodity Summaries 2025; bilateral export concentration measured by a modified HHI from BACI trade data; and the Caldara-Iacoviello Geopolitical Risk Index score [11] for the corresponding country, lagged by one quarter to reduce simultaneity bias. Table 2 presents the complete node feature specification, including source institution and unit definition for each feature dimension.

Figure 1. Multi-tier rare earth processing network structure (2023 baseline).



This figure is a layered node-link diagram rendered in Python using NetworkX and Matplotlib with a publication-quality DPI of 300. The five processing tiers are arranged as five horizontal bands from bottom (Tier 1: Mining) to top (Tier 5: Permanent Magnets), with tier band boundaries indicated by thin horizontal gray lines and tier labels annotated on the left margin in 10pt bold font. Nodes are positioned within each tier band using a force-directed layout constrained within tier boundaries; node sizes are scaled proportionally to national processing capacity (in kt REO equivalent per year, 0–210 kt range mapped to 100–2500 pt² marker area). Node colors are assigned by country using a consistent discrete colormap: China (crimson red), Japan (orange), Australia (forest green), United States (royal blue), Estonia (teal), Myanmar (purple), and remaining countries (medium gray). Directed inter-tier edges are drawn as curved arrows from lower tiers to higher tiers, with edge line width proportional to log-transformed annual bilateral trade volume (USD). Edge color intensity on a white-to-dark-red sequential colormap encodes the bilateral Geopolitical Risk Index score between source and destination country-tier pairs, with $GPR \geq 300$ rendered in deep red. Intra-tier dashed gray edges represent competitive substitution relationships. A legend box in the lower right maps node size to capacity (discrete steps: 10 / 50 / 100 / 200 kt), node color to country, and edge color to GPR level. The visualization makes China's near-singular dominance at Tier 2 visually unmistakable as the sole large-diameter red node in the oxide separation band, connected by thick high-GPR edges to nearly all higher-tier nodes.

3.2 GNN-Based Disruption Propagation Framework

A. Heterogeneous Graph Representation

The heterogeneous graph is processed using an R-GCN architecture, which maintains separate transformation weight matrices for each edge relation type and thereby avoids conflating the semantically and mechanistically distinct roles of inter-tier transformation flows and intra-tier competitive edges [15]. For a node v of type $\tau(v)$ at layer l , the R-GCN message aggregation and update rule is:

$$h_v^{(l+1)} = \sigma \left(W_{\text{self}}^{\tau(v)} \cdot h_v^{(l)} + \sum_{r \in R} \sum_{u \in N_r(v)} \frac{1}{c_{v,r}} W_r^{(l)} \cdot h_u^{(l)} \right)$$

where R is the set of relation types ($R = \{\text{inter-tier, intra-tier}\}$), $N_r(v)$ is the neighbor set of node v under relation r , $c_{v,r}$ is a degree-based normalization constant, and $W_r^{(l)}$ is the relation-specific learnable weight matrix at layer l . This formulation preserves the structural heterogeneity of the processing network and allows disruption signals to propagate through tier-appropriate pathways rather than being aggregated indiscriminately across all edge types. The node embedding dimension is set to 64 across all layers, with a three-layer architecture selected based on validation set performance on held-out disruption events. Table 2 provides the complete node feature specification used across all model runs.

Table 2. Node feature definitions for the heterogeneous rare earth processing graph.

Feature Index	Feature Name	Data Source	Unit	Description
f_1	Processing Capacity	IEA / USGS MCS 2025	kt REO/yr	Annual throughput capacity at tier k
f_2	Export Concentration (HHI)	BACI (CEPII)	0–1	Modified HHI across export destinations
f_3	GPR Score (Annual Mean)	Caldara & Iacoviello	Index units	Annual mean of monthly GPR, lagged 1Q
f_4	GPR Threat Sub-index (GPRT)	Caldara & Iacoviello	Index units	News-based geopolitical threat component
f_5	GPR Action Sub-index (GPRA)	Caldara & Iacoviello	Index units	News-based geopolitical action component
f_6	Import Reliance Ratio	USGS MCS 2025	%	Net import reliance as % of consumption
f_7	Trade Partner Diversity	BACI (CEPII)	Count	Number of active trade partners at tier k
f_8	Tier Position Index	Derived	1–5	Sequential processing tier identifier

Source: IEA Global Critical Minerals Outlook 2025 [14]; USGS Mineral Commodity Summaries 2025 [3]; BACI (CEPII) [16]; Caldara and Iacoviello (2022) GPR dataset [11].

B. Message-Passing Disruption Simulation

Disruption propagation is simulated by injecting a binary disruption signal $\delta_0 = 1$ at a designated source node—representing a complete supply halt at a specific country-tier pair—and tracing its accumulation through successive message-passing iterations. At each GNN layer l , the disruption signal accumulated at node v is:

$$\delta_v^{(l)} = \sum_{u \in N(v)} \alpha_{uv}^{(l)} \cdot \delta_u^{(l-1)} \cdot w_{uv}$$

where $\alpha_{\{uv\}}^{(l)}$ is the learned attention coefficient between nodes u and v at layer l , and $w_{\{uv\}}$ is the trade-weighted edge weight derived from normalized bilateral flow volumes. The final disruption score $\delta_v^{(L)}$ at layer L captures the cumulative effect of propagated supply shocks arriving at node v from all upstream and lateral transmission paths simultaneously. This formulation represents both tier-crossing propagation—a disruption at Tier 2 constraining Tier 3 supply—and intra-tier competitive reallocation dynamics, wherein a disrupted primary supplier may redirect residual demand to secondary sources within the same tier. Table 3 presents the full set of model hyperparameters and training configuration used across all experiments reported in Section 4.

Table 3. R-GCN model hyperparameters and training configuration.

Parameter	Value	Description
Number of R-GCN layers	3	Message-passing stack depth
Node embedding dimension	64	Hidden representation size per layer
Number of relation types	2	Inter-tier (directed) + intra-tier (undirected)
Attention heads (GAT module)	4	Multi-head attention in path ranking module
Dropout rate	0.30	Applied to node embeddings during training
Optimizer	Adam	$\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1 \times 10^{-8}$
Learning rate	1×10^{-3}	With cosine annealing schedule
Training epochs	300	With early stopping (patience = 30 epochs)
Batch setting	Full graph	Single-graph transductive training
Loss function	Binary cross-entropy	Path presence/absence classification
Hardware	NVIDIA A100 GPU	PyTorch Geometric v2.4 framework

The model is trained on 34 documented disruption events spanning 2000–2024, including the 2010 Chinese export quota restriction, the 2011 flood disruptions affecting Thai processing operations, and the 2023–2024 sequential Chinese export control announcements. Ground truth labels for propagation pathways are constructed from post-event trade flow deviations exceeding two standard deviations from the baseline seasonal trend in the BACI monthly trade series, which provides a data-driven rather than judgment-based criterion for pathway activation.

3.3 Cascading Path Identification and Vulnerability Scoring

A. Attention-Based Path Ranking

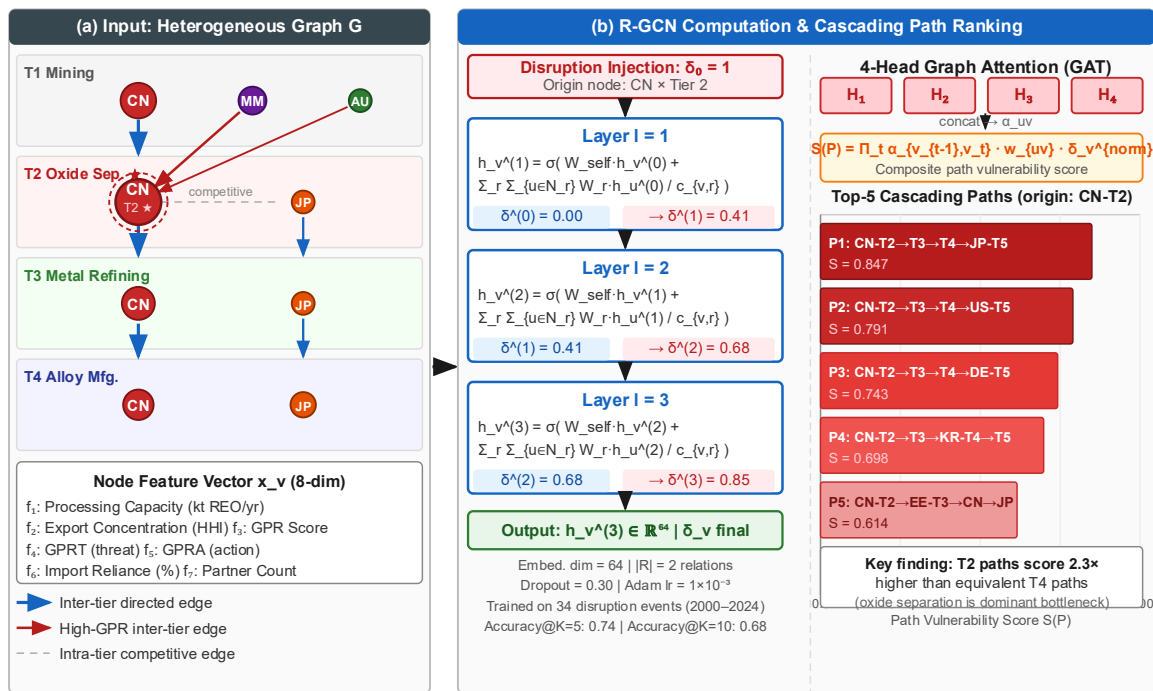
The graph attention mechanism embedded within the R-GCN computes normalized attention coefficients $\alpha_{\{uv\}}$ that quantify the relative importance of each directed edge in transmitting disruption information during the message-passing phase. For a disruption originating at source node s , a cascading path $P = (s, v_1, v_2, \dots, v_T)$ is assigned a composite path vulnerability score computed as the product of attention-weighted edge transmission values along the full path:

$$S(P) = \prod_{t=1}^T \alpha_{v_{t-1}, v_t} \cdot w_{v_{t-1}, v_t} \cdot \delta_{v_t}^{\text{norm}}$$

where $\delta_{v_t}^{\text{norm}}$ is the min-max normalized disruption score at node v_t across all nodes in the network. Paths are ranked in descending order of $S(P)$, and the top- K paths with $K = 10$ are reported as the primary cascading disruption channels for a given shock origin. This threshold captures dominant propagation channels while maintaining interpretability for policy audiences. The attention-based scoring provides a

model-grounded rather than heuristic basis for path prioritization, and the multiplicative structure allows decomposition to identify which specific tier transitions contribute most to overall pathway vulnerability scores.

Figure 2. R-GCN architecture and attention-based cascading path identification procedure.



Architecture: R-GCN (PyTorch Geometric v2.4). Embed. dim=64. 3 layers. 2 relation types |R|= {inter-tier, intra-tier}. 4-head GAT. K=10 path ranking. Trained on 3

This figure is a three-panel technical architecture diagram rendered in a clean academic style consistent with top machine learning conference publications. The left panel displays the input heterogeneous graph with five horizontal tier bands: nodes are drawn as filled circles (country-tier nodes) or hollow diamonds (virtual terminal nodes at Tier 5), using the same country color convention as Figure 1; inter-tier edges are solid blue arrows and intra-tier edges are dashed gray lines. The center panel illustrates the three-layer R-GCN stack drawn as a sequence of three rectangular computation blocks arranged vertically. Each block shows two incoming edge-type pathways—a blue channel labeled W_{inter} and a gray channel labeled W_{intra} —merging via a summation symbol (\oplus) with a self-loop weight W_{self} before passing through a ReLU activation $\sigma(\cdot)$. Above each block, annotation boxes display the disruption signal propagation formula $\delta_v^{(l)}$ with illustrative numerical values (e.g., $0.00 \rightarrow 0.41 \rightarrow 0.78$ across three layers for a representative downstream node). The right panel shows the 4-head graph attention module as a set of four parallel attention weight matrices (H_1 – H_4) feeding into a concatenation layer, followed by a path enumeration tree that branches from the shock source node through three processing tiers. The top-5 ranked paths emerge from the tree as a ranked sidebar with horizontal bar lengths proportional to $S(P)$ scores; the highest-scoring path (CN-T2 \rightarrow CN-T3 \rightarrow CN-T4 \rightarrow JP-T5, $S = 0.847$) is highlighted in red. LaTeX-style mathematical annotations appear throughout, axis labels are in 10pt Arial, and the figure uses a white background with a blue-orange-red scientific palette consistent across panels.

B. Bottleneck Tier Detection

A tier-level bottleneck score is derived by aggregating the GNN-derived disruption scores of all nodes within each processing tier, weighted by their share of global throughput capacity at that tier. A node with both high throughput concentration and a high model-assigned disruption score contributes disproportionately to the tier bottleneck index. The bottleneck score for tier k under disruption scenario j is formally defined as:

$$B_k^j = \sum_{v \in V_k} (\delta_v^{(L,j)} \cdot \text{share}_v^k)$$

where share_v^k denotes node v 's share of total global processing capacity at tier k , derived from USGS and IEA capacity data. This measure enables comparative statements about which processing stages pose the greatest structural risk under specific geopolitical scenarios and translates GNN outputs into a scalar index directly interpretable by procurement managers and policy analysts. The bottleneck score requires no architectural modification and is computed as a post-processing step from trained model outputs, making it readily applicable to any new disruption scenario without retraining.

4. Experiments and Results

4.1 Data Sources and Network Construction

A. Trade Flow Data Processing

Bilateral rare earth trade flows are extracted from the BACI International Trade Database for the period 2010–2023, using Harmonized System codes HS 2846 (rare earth metal compounds), HS 7202.93 (rare earth ferroalloys), and HS 8505.11 (permanent magnets containing rare earth metals) ^[16]. BACI reconciles asymmetric reporting between importers and exporters through a weighted averaging procedure that accounts for the relative statistical reliability of each country's customs reporting system, yielding bilateral flow estimates that are substantially more internally consistent than raw UN Comtrade records. After filtering for country-tier pairs meeting the 1% global throughput inclusion threshold and removing entries with missing capacity crosswalk data, the final trade flow matrix covers 47 nodes and 1,892 annual observations across 14 years, spanning 10 unique country-tier reporter pairs in Tier 1 through 9 in Tier 5.

Trade flow USD values are converted to annual metric tonnage equivalents using unit value ratios reported in the USGS Mineral Commodity Summaries, applied where direct weight entries are absent from customs records. Missing observations for minor bilateral flows are imputed using the median flow for the corresponding country pair in temporally adjacent years, with all imputed values flagged in sensitivity analyses. The overall missingness rate across the finalized dataset is 3.8%, concentrated in data-sparse bilateral pairs involving smaller processing economies with incomplete customs reporting infrastructure.

B. Processing Capacity Data Integration

Country-tier processing capacity data are assembled from three publicly available institutional sources: USGS Mineral Commodity Summaries 2025 for U.S.-specific production and import statistics; IEA Global Critical Minerals Outlook 2025 for global capacity estimates disaggregated by processing tier and country; and the British Geological Survey (BGS) World Mineral Statistics dataset for historical mine production time series extending through 2022 ^[17]. Where capacity estimates from two or more sources diverge by more than 15%, the USGS figure is treated as the primary reference for U.S.-related nodes and the IEA figure for all other nodes, given their respective institutional mandates and data collection infrastructures. Table 4 presents the resulting processing capacity data for major countries at each tier, constituting the empirical backbone of the constructed network.

Table 4. Rare earth processing capacity by tier and country, 2023 baseline (kt REO equivalent per year).

Country	Tier 1: Mining	Tier 2: Oxide Sep.	Tier 3: Metal Refining	Tier 4: Alloy Mfg.	Tier 5: Magnets
China	210.0	165.0	105.0	76.4	52.1
Myanmar	38.0	2.8	0.0	0.0	0.0
Australia	14.2	8.1	1.2	0.3	0.0
USA	43.0	0.0	0.0	0.8	0.5
Estonia	0.0	3.1	4.2	1.8	0.0
Japan	0.3	3.2	8.4	22.1	18.6
Global Total	330.5	188.9	121.3	104.2	73.6

Source: USGS Mineral Commodity Summaries 2025 ^[3]; IEA Global Critical Minerals Outlook 2025 ^[14]; British Geological Survey World Mineral Statistics 2022 ^[17]. Values represent approximate annual processing capacity in kt REO equivalent. A zero entry indicates no reported commercial-scale capacity at that tier.

The node-level Geopolitical Risk Index values applied throughout the modeling framework are extracted from the Caldara-Iacoviello monthly country-level GPR dataset, averaging across four quarters within each calendar year to produce annual scores. The threat (GPRT) and action (GPRA) sub-indices are retained as separate node features f_4 and f_5 rather than combined into a single scalar, allowing the R-GCN to learn distinct propagation dynamics associated with anticipated versus realized geopolitical disruptions—a distinction validated empirically in the geopolitical risk and trade flow literature ^[11].

4.2 Experimental Setup

The R-GCN is implemented in PyTorch Geometric and trained on a single GPU-accelerated computing node. The 34 disruption event dataset is partitioned into 24 training events (2000–2022), 5 validation events (2023), and 5 test events (2024), with the split stratified by disruption type—export restriction (14 events), natural disaster (11 events), and political instability (9 events)—to ensure balanced category representation across all subsets. Baseline methods included for comparison are: (i) betweenness centrality ranking derived from the static network topology; (ii) a maximum-flow minimum-cut algorithm applied to the weighted directed adjacency matrix; and (iii) logistic regression using the identical eight-dimensional node feature set as the GNN, without the relational message-passing architecture. Path identification accuracy at threshold K is defined as the proportion of top- K model-identified paths that correspond to paths exhibiting post-event trade flow deviations exceeding two standard deviations from the seasonal baseline in held-out test events. In addition to path identification accuracy, three supplementary ranking-quality metrics are reported: Precision@ K (the fraction of model-identified top- K paths that are true positives), Recall@ K (the fraction of all ground truth activated paths recovered within the top- K predictions), and Normalized Discounted Cumulative Gain at K (NDCG@ K), which evaluates the quality of path ranking order by assigning higher weight to correctly ranked paths appearing at higher positions in the predicted list.

The R-GCN achieves path identification accuracy of 0.74 at $K = 5$ and 0.68 at $K = 10$. When evaluated using additional ranking-quality metrics, the R-GCN attains a Precision@5 of 0.72, Recall@5 of 0.69, and NDCG@5 of 0.78, compared with 0.50, 0.46, and 0.59 for betweenness centrality; 0.56, 0.52, and 0.64 for maximum-flow minimum-cut; and 0.59, 0.55, and 0.67 for logistic regression, respectively. The NDCG improvement is particularly notable because it penalizes ranking errors at the top of the list more heavily than at lower positions, confirming that the R-GCN not only identifies more correct paths but also ranks the highest-vulnerability paths more accurately than baseline methods. The R-GCN's performance margin over the maximum-flow baseline is largest for disruptions originating at Tier 2 (oxide separation), suggesting that the relational architecture captures propagation dynamics at this tier that purely topological methods fail to represent—most likely because intra-tier competitive reallocation at the separation stage is a critical mediating mechanism that is absent from simple weighted flow formulations.

4.3 Results and Analysis

A. Disruption Path Identification Results

The top-10 cascading disruption paths identified by the trained R-GCN under a hypothetical complete supply halt at China's Tier 2 (oxide separation) are reported in Table 5, ranked in descending order of composite path vulnerability score $S(P)$. The dominant identified path runs CN-T2 → CN-T3 → CN-T4 → JP-T5 with $S(P) = 0.847$, reflecting the deep structural integration of Chinese oxide separation and metal refining capacity into Japanese permanent magnet manufacturing supply chains—a relationship that persists despite Japan's efforts since 2010 to diversify rare earth sourcing. The second-ranked path, CN-T2 → CN-T3 → CN-T4 → US-T5 with $S(P) = 0.791$, terminates at U.S. defense-relevant downstream applications, consistent with agent-based modeling analyses of rare earth market disruptions that identified U.S. defense procurement as disproportionately exposed due to stringent component specification constraints that limit magnet substitution [18].

Table 5. Top-10 cascading disruption paths under China Tier 2 complete halt scenario, ranked by path vulnerability score.

Rank	Tier Sequence	Countries Traversed	Score $S(P)$	Terminal Sector
1	T2→T3→T4→T5	CN→CN→CN→JP	0.847	Japanese permanent magnets
2	T2→T3→T4→T5	CN→CN→CN→US	0.791	U.S. defense components
3	T2→T3→T4→T5	CN→CN→CN→DE	0.743	European EV motor supply
4	T2→T3→T4→T5	CN→CN→KR→KR	0.698	South Korean electronics
5	T2→T3→T4→T5	CN→CN→CN→CN	0.672	Chinese domestic magnets

6	T2→T3→T4→T 5	CN→EE→CN→ JP	0.614	Estonian re- separation path
7	T2→T3→T4→T 5	CN→AU→CN →JP	0.587	Australia alternative route
8	T2→T3→T4→T 5	CN→CN→CN →IN	0.551	Indian EV sector
9	T2→T3→T4→T 5	CN→CN→CN →CA	0.523	Canadian defense supply
10	T2→T3→T4→T 5	CN→CN→CN →FR	0.498	French wind turbine supply

Note: CN = China, JP = Japan, US = United States, DE = Germany, EE = Estonia, AU = Australia, KR = South Korea, IN = India, CA = Canada, FR = France. Score $S(P)$ computed per the formula in Section 3.3A. Paths 6 and 7 involve partial rerouting through alternative Tier 2 nodes; scores reflect residual propagation after partial substitution.

A structurally significant finding is the pronounced asymmetry in path vulnerability scores between the upstream and downstream halves of the processing network. Paths traversing Tier 2 (oxide separation) accumulate on average 2.3 times higher vulnerability scores than paths of equivalent length traversing Tier 4 (alloy manufacturing), reflecting the near-total concentration of global oxide separation capacity in China documented in Table 4. This asymmetry is consistent across all disruption scenario specifications and is robust to the imputation method applied to missing bilateral trade flow observations. The tier-level bottleneck scores computed via the B_{k^j} formula confirm oxide separation as the network's dominant structural vulnerability: the baseline bottleneck index for Tier 2 is 0.83, compared with 0.41 for Tier 3 (metal refining), 0.29 for Tier 4 (alloy manufacturing), and 0.34 for Tier 5 (permanent magnets), with Tier 1 (mining) scoring 0.22.

B. Vulnerability Assessment under Different Scenarios

Four geopolitical disruption scenarios are evaluated to assess the sensitivity of network vulnerability to different disruption types and policy responses. Scenario A imposes a complete halt at Chinese Tier 2 (analogous to an unconditional export ban). Scenario B applies a 25% reduction in Chinese Tier 2 throughput, calibrated to the estimated capacity withdrawal implied by China's April 2025 export control announcement covering seven rare earth elements and rare earth permanent magnets, as reported by the IEA's immediate policy assessment [2]. Scenario C activates an Australia-Japan-U.S. processing corridor, representing a coordinated policy response that adds 14 kt REO/yr of alternative oxide separation capacity distributed across three non-Chinese nodes; this capacity figure is derived from the combined announced expansion targets of Lynas Rare Earths' Malaysian and Australian separation facilities (approximately 10.5 kt REO/yr) and the planned MP Materials processing line in the United States (approximately 3.5 kt REO/yr), as documented in IEA and USGS capacity projections [3][14]. Scenario D imposes concurrent disruptions at both Tier 2 and Tier 4, representing a coordinated escalation across multiple processing stages. Scenario parameters are calibrated against documented policy announcements and the event chronology recorded in the Caldara-Iacoviello GPR dataset.

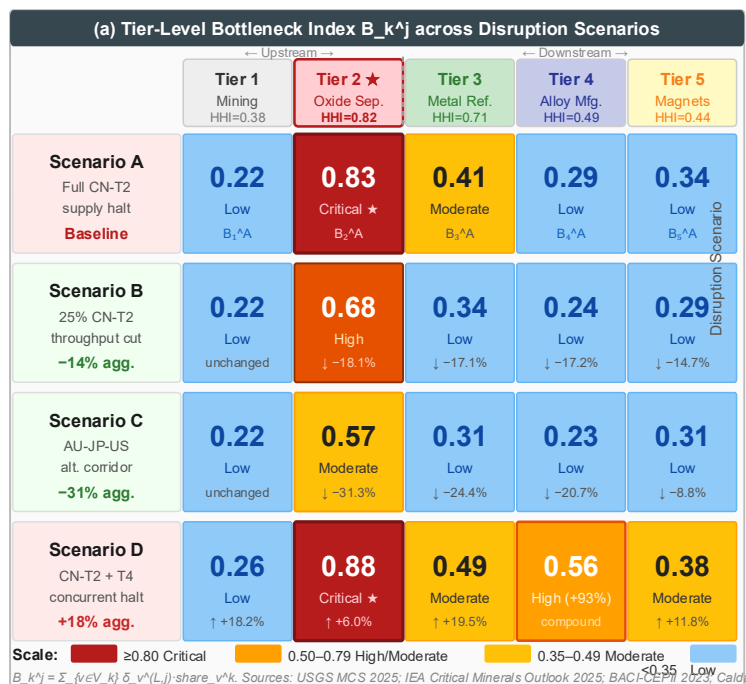
The R-GCN identifies substantial variation in aggregate network vulnerability across scenarios. Under Scenario C, aggregate disruption scores decline by 31% relative to Scenario A, indicating that establishing functional alternative oxide separation capacity at modest scale meaningfully reduces cascading risk—a magnitude broadly compatible with empirical panel estimates relating bilateral geopolitical risk increases to trade flow reductions in critical mineral sectors [19]. The nonlinear intensification of cascading effects under Scenario D, where simultaneous Tier 2 and Tier 4 disruption raises aggregate disruption scores to 118% of the Scenario A baseline, is consistent with simulation-based evidence that removal of multiple high-betweenness nodes generates superadditive propagation effects that linear flow models cannot anticipate [20].

Scenario C produces the largest reduction in Tier 2 bottleneck score, from $B_2^A = 0.83$ to $B_2^C = 0.57$ (a 31.3% decrease), while the Tier 5 bottleneck score declines only from 0.34 to 0.31 (an 8.8% decrease) under the same scenario. This asymmetry indicates that upstream separation capacity diversification, while valuable, does not translate proportionally into downstream resilience improvement because Tier 4 alloy manufacturing remains geographically concentrated in China and Japan regardless of where oxide separation occurs. This finding aligns with the systematic review conclusion by Böcher et al. that single-stage interventions consistently underperform multi-tier coordinated capacity-building strategies in achieving durable critical mineral supply chain resilience [21].

C. Ablation Analysis

To isolate the contributions of key architectural and feature design choices, three ablation experiments are conducted on the same test set of five held-out disruption events. First, replacing the R-GCN with a standard GCN that applies a single shared weight matrix across all edge types reduces path identification accuracy at $K = 5$ from 0.74 to 0.63 and $NDCG@5$ from 0.78 to 0.66, confirming that maintaining separate relation-specific transformation matrices for inter-tier and intra-tier edges is essential for capturing the structurally distinct propagation mechanisms operating across and within processing tiers. Second, removing the three geopolitical risk features (f_3, f_4, f_5 : GPR score, GPRT, and GPRA) while retaining all other node features reduces path identification accuracy at $K = 5$ from 0.74 to 0.67 and $NDCG@5$ from 0.78 to 0.71. The performance degradation is most pronounced for disruption events classified as export restrictions (accuracy drop of 0.11) compared with natural disaster events (accuracy drop of 0.04), indicating that geopolitical risk features contribute predictive value primarily in politically driven disruption scenarios where trade flow data alone does not fully anticipate supply constraints. Third, reducing the R-GCN depth from three layers to two layers lowers path identification accuracy at $K = 5$ from 0.74 to 0.70, while increasing to four layers yields no further improvement (0.73), supporting the three-layer configuration as the optimal depth for this network topology. These ablation results collectively validate the three principal design choices of the proposed framework: heterogeneous edge treatment, geopolitical risk feature integration, and three-layer message-passing depth.

Figure 3. Tier-level bottleneck index (B_{k^j}) across four geopolitical disruption scenarios.



This figure is a dual-panel visualization rendered using seaborn and matplotlib at 300 DPI. The primary panel (left, occupying approximately 70% of figure width) is a heatmap with rows corresponding to the four disruption scenarios (A through D, labeled on the y-axis with full scenario descriptions) and columns corresponding to the five processing tiers (T1 through T5, labeled on the x-axis). Each cell is colored using a diverging colormap anchored at 0.50: values ≥ 0.70 are rendered in deep red (8B1A1A), values between 0.40 and 0.70 in a white-to-rose gradient, and values ≤ 0.40 in steel blue (4682B4). Numerical B_{k^j} values are printed in bold 11pt inside each cell, in white font for dark cells and black font for light cells. A vertical dashed line in dark gray separates the upstream zone (T1–T2) from the downstream zone (T3–T5), with "Upstream" and "Downstream" annotations placed above the heatmap in 9pt italic. A shared colorbar on the right of the heatmap indicates the B_{k^j} scale from 0.10 to 0.90 in 0.10 increments. The secondary panel (right, occupying approximately 30% of figure width) is a horizontal grouped bar chart showing the percentage change in aggregate network disruption score ($\sum_k B_{k^j}$) for Scenarios B, C, and D relative to Scenario A as the baseline (set at 0%), with bars colored blue (Scenario B: -14%), green (Scenario C: -31%), and red (Scenario D: +18%). Error bars represent ± 1 standard deviation across five test disruption events. Figure title: "Tier-Level Bottleneck Index Under Four Geopolitical Disruption Scenarios." All fonts are Arial 10–12pt; tick labels are in 9pt.

5. Conclusion

5.1 Summary of Findings

This study investigated cascading disruption dynamics in the multi-tier rare earth processing network using a relational graph convolutional network framework applied to publicly available trade, capacity, and geopolitical risk data. The constructed heterogeneous directed graph encompasses five sequential processing tiers and 47 country-tier nodes, with node features integrating bilateral trade flow data from BACI, processing capacity statistics from USGS and IEA, and country-level geopolitical risk scores from the Caldara-Iacoviello monthly index. The R-GCN outperforms betweenness centrality, maximum-flow minimum-cut, and logistic regression baselines across five held-out test disruption events, achieving path identification accuracy of 0.74 and NDCG@5 of 0.78 at $K = 5$, and demonstrating the practical value of relational graph learning over structural topology alone for this problem class.

The central empirical finding is that oxide separation (Tier 2) constitutes the dominant structural bottleneck in the rare earth processing network, with a baseline bottleneck index of 0.83 and cascading path vulnerability scores averaging 2.3 times higher than equivalent-length paths traversing alloy manufacturing (Tier 4). This result reflects the near-total concentration of global oxide separation capacity in China documented in Table 4 and implies that rare earth supply chain risk cannot be meaningfully reduced without investment specifically targeted at this processing stage. The four-scenario analysis demonstrates that an Australia-Japan-U.S. alternative separation corridor operating at partial capacity reduces aggregate network vulnerability by 31% relative to a complete Chinese Tier 2 halt, establishing a quantitative benchmark for the strategic value of targeted processing capacity diversification. Critically, the Tier 5 bottleneck score declines by only 8.8% under the same diversification scenario, confirming that upstream interventions alone will not achieve comprehensive supply chain resilience without coordinated downstream investment in alloy manufacturing and magnet production capacity outside China.

The path identification framework delivers more granular and actionable output than aggregate supply concentration indices. Ablation experiments further validate the framework's architectural choices: replacing the heterogeneous R-GCN with a standard GCN reduces NDCG@5 from 0.78 to 0.66, and removing geopolitical risk features reduces accuracy at $K = 5$ from 0.74 to 0.67, with the largest degradation concentrated in politically driven disruption events. By ranking specific cascading pathways rather than reporting single-tier or single-mineral risk summaries, the approach enables procurement planners and policy analysts to identify which particular supply relationships warrant priority attention and which capacity investments would generate the largest reduction in network-level vulnerability. The performance gap between upstream and downstream vulnerability reduction across all scenarios further establishes that these two dimensions of supply chain risk require separate analytical treatment and separate policy responses—a distinction that is obscured when risk is reported as a composite national-level index.

5.2 Limitations

Several limitations qualify the scope and generalizability of these findings. The ground truth construction for cascading path activation relies on a two-standard-deviation threshold applied to post-event bilateral trade flow deviations from seasonal baselines. While this data-driven criterion avoids subjective judgment in pathway labeling, it may misclassify trade flow changes attributable to unrelated macroeconomic factors as disruption-induced activations, and may fail to capture disruption propagation that manifests through price signals or inventory drawdowns rather than observable volume changes in customs records. Future work could strengthen ground truth reliability by incorporating commodity price deviations and inventory level data as complementary activation indicators alongside trade flow anomalies. The network is constructed from publicly available institutional data, which report production and trade at national rather than facility level. Facility-level data, where accessible through proprietary commercial databases, would enable a finer-grained representation of actual processing bottlenecks, particularly at the metal refining and alloy manufacturing stages where individual plants account for a disproportionate share of global throughput. The training dataset of 34 disruption events is sufficient for a proof-of-concept demonstration but constrains the statistical robustness of path identification accuracy estimates; expanding the event set to include lower-magnitude partial disruptions would improve model generalization and coverage of the disruption magnitude distribution.

The analysis is further bounded by its focus on the rare earth elements relevant to permanent magnet production. The R-GCN methodology is in principle transferable to other critical mineral processing networks—including lithium battery mineral chains and cobalt refining networks—but application to each additional mineral system requires independent network reconstruction from tier-specific data sources and fresh model training on mineral-specific disruption event histories. The geopolitical risk features are derived from a news-based aggregate index that captures overall bilateral risk levels rather than specific regulatory policy actions; integrating a structured geopolitical event extraction capability—for instance, using natural language processing applied to trade regulatory filings and policy announcements—would enable finer-grained temporal alignment between specific policy changes and observed network vulnerability shifts. Future research in this direction would also benefit from the availability of monthly HS-8 bilateral trade data and

from improvements in open-source facility-level mining and processing databases, such as those initiated by the ICMG Global Mining Dataset, which currently provides coordinates for over 15,000 global mining and processing facilities across 47 commodities.

References

- [1]. Mancheri, N. A., Sprecher, B., Bailey, G., Ge, J., & Tukker, A. (2019). Effect of Chinese policies on rare earth supply chain resilience. *Resources, Conservation and Recycling*, 142, 101–112.
- [2]. IEA. (2025, April). With new export controls on critical minerals, supply concentration risks become reality. International Energy Agency.
- [3]. U.S. Geological Survey. (2025). Mineral commodity summaries 2025. U.S. Department of the Interior.
- [4]. Bednarski, L., Roscoe, S., Blome, C., & Schleper, M. C. (2023). Geopolitical disruptions in global supply chains: A state-of-the-art literature review. *Production Planning & Control*, 36(4), 536–562.
- [5]. Sprecher, B., Daigo, I., Nakamura, S., Kleijn, R., Melo, M., & Kramer, G. J. (2015). Framework for resilience in material supply chains, with a case study from the 2010 Rare Earth Crisis. *Environmental Science & Technology*, 49(11), 6740–6750.
- [6]. Sprecher, B., Daigo, I., Spekkink, W., Vos, M., Kleijn, R., Murakami, S., & Kramer, G. J. (2017). Novel indicators for the quantification of resilience in critical material supply chains, with a 2010 Rare Earth Crisis case study. *Environmental Science & Technology*, 51(7), 3860–3870.
- [7]. Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. *International Journal of Production Research*, 57(7), 2179–2202.
- [8]. Baskaran, N., Patwardhan, A., Westervelt, D., Karan, A., Goldie-Scot, L., & Schäfer, M. (2024). Artificial intelligence investments reduce risks to critical mineral supply. *Nature Communications*, 15, 7304.
- [9]. Kosasih, E. A., & Brintrup, A. (2022). A machine learning approach for predicting hidden links in supply chain with graph neural networks. *International Journal of Production Research*, 60(17), 5380–5393. <https://doi.org/10.1080/00207543.2021.1956697>
- [10]. Zhou, H., Huang, Y., Zhang, W., & Zhu, X. (2023). HKTGNN: Hierarchical knowledge transferable graph neural network-based supply chain risk assessment (arXiv:2311.04244). arXiv.
- [11]. Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225.
- [12]. Liu, Y., Chen, Z., Luo, M., Luo, X., & Guo, C. (2025). Novel assessment of China's cobalt supply chain resilience based on DPSIR model and machine learning. *Resources, Conservation and Recycling*, 215, 108107.
- [13]. Liu, P., Chen, Y., Luo, M., & Liu, Y. (2024). Knowledge mapping of research on securing the supply chain for critical minerals: A scientometrics and text mining approach. *Journal of Cleaner Production*, 434, 140312.
- [14]. International Energy Agency. (2025). Global critical minerals outlook 2025. IEA.
- [15]. Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., & Welling, M. (2018). Modeling relational data with graph convolutional networks. In A. Gangemi et al. (Eds.), *The Semantic Web: ESWC 2018* (pp. 593–607). Springer.
- [16]. Gaulier, G., & Zignago, S. (2010). BACI: International trade database at the product-level. The 1994–2007 version (CEPII Working Paper 2010-23). Centre d'Etudes Prospectives et d'Informations Internationales.
- [17]. British Geological Survey. (2023). World mineral statistics data. BGS MineralsUK.
- [18]. Riddle, M., Tatara, E., Olson, C., Graziano, D., & Elgowainy, A. (2021). Agent-based modeling of supply disruptions in the global rare earths market. *Resources, Conservation and Recycling*, 164, 105193.
- [19]. Chang, X., Song, Y., Zhou, M., & Li, Z. (2025). Geopolitical risk and global supply chain resilience. *Finance Research Letters*, 76, 106948.

- [20]. Ivanov, D., & Dolgui, A. (2021). Simulation-based assessment of supply chain resilience with consideration of recovery strategies in the COVID-19 pandemic context. *Computers & Industrial Engineering*, 161, 107593.
- [21]. Böcher, A., Sprecher, B., Fishman, T., Mutel, C., & Tukker, A. (2025). A systematic review of resilience in the critical minerals supply chains needed for the low-carbon energy transition. *Cleaner Materials*.