

A GNN-Based Framework for Multi-Objective Telecommunications Quality Ranking

Gabriel Luiz Espindola Pedro, Paulo Ricardo Branco da Silva, Dimas Irion Alves, João Vitor Rosa Negri, Arthur Cadore Matuella Barcella, Renato Machado, Kai de Andrade Lima, and Paulo Roberto Silva Marcos

Abstract—Assessing telecommunications mobile services quality across municipalities with heterogeneous structural, geographic, and socioeconomic conditions requires methodologies beyond simple scalar or functional aggregations. This paper introduces a self-supervised Graph Neural Network (GNN) based ranking framework that learns the geometry of municipality-level Quality of Service (QoS) indicators on a k -nearest-neighbor graph, producing a unified ranking obtained by self-supervised learning on Pareto-inspired proxies. The architecture employs residual graph convolutional layers and is trained with a multi-objective loss function that enforces pairwise ranking consistency, input reconstruction fidelity, geometric smoothness on the learned manifold, and directional alignment with a Pareto-inspired proxy. The approach is validated on a dataset covering Brazilian municipalities over multiple monthly periods, totaling over 82,000 municipality-month observations. Results show that the proposed model produces scores with a wider effective range ($\sigma = 0.098$) than ANATEL's legacy QoS ($\sigma = 0.054$) and its revised version ($\sigma = 0.058$) baselines, improving rank discrimination. Distribution and spatial analyses suggest that the framework produces more informative and distinguishable municipality rankings than the baselines, mitigating the concentration effects observed in current regulatory composite indicators.

Keywords—Learn to Rank, Graph Neural Networks, Pareto Dominance, Multi-Objective Optimization, Quality of Service

I. INTRODUCTION

International best practices advocate results-based regulatory approaches for Quality of Service (QoS). Given the multidimensional nature of QoS indicators, municipality ranking provides support in the prioritization of regulatory actions. Recent works have explored data-driven approaches for ranking that go beyond traditional aggregation, including graph-neural-network-based QoS prediction [1], [2], deep-learning-based Quality of Experience (QoE) assessment for wireless networks [3], and modern Learning to Rank techniques that learn scoring functions directly from data [4].

Defining a suitable ranking strategy is not trivial. The multidimensional nature of telecommunications quality—encompassing network and customer-related aspects—means

that different indicators can lead to conflicting orderings, a characteristic of multi-objective optimization problems [5], [6]. Establishing a definitive ground-truth ranking is difficult; evaluating model outputs relies on partial-ordering frameworks and domain-specific contextual knowledge. Selecting the essential variables for ranking definition is equally unclear: not only QoS indicators are required, but also covariates such as infrastructure and QoE indicators.

In multi-objective settings, Pareto dominance provides a natural partial ordering: municipality A dominates municipality B if A is at least as good in every dimension and strictly better in at least one [7]. However, many municipality pairs are not comparable under strict Pareto ordering, motivating the search for a principled scalarization rule that extends this partial order into a total order while still respecting the dominance structure of the data. The core challenge is therefore not only aggregation, but learning a scoring function that faithfully accounts for each attribute's contribution—so that entities which differ in even one dimension are consistently and meaningfully separated in the induced ranking.

Brazilian National Telecommunications Agency (ANATEL) consolidates QoS indicators at the municipal level through the Service Quality Index (IQS), a linear composite based on weighted aggregation for each provider and semester [8], [9]. This approach aligns with international practices and reflects a broader trend of deriving orderings from multidimensional QoS metrics via linear aggregation. However, such non-relational formulations may fail to capture the global structure of relative municipal performance. By ignoring Pareto dominance relationships, scores become less discriminative, with similar entities assigned near-identical values. As detailed in Section IV, the legacy IQS concentrates 50% of observations between 0.835 and 0.863, an interquartile range of only 0.028.

To address these challenges, we propose a Graph Neural Networks (GNN)-based Pareto-inspired model that learns, through pairwise competition between entities, a scalar quality score $\hat{r}_i$ for each municipality i . A k -Nearest Neighbors (k NN) graph over the feature space provides the relational structure for this learning process. Experimental results indicate that the proposed multi-objective scheme expands score dispersion and improves rank discriminability, yielding a more favorable optimization manifold than unidimensional scalar approaches.

The remainder of the paper is organized as follows. Section II describes the dataset of municipality-level QoS, QoE, and infrastructure indicators provided by ANATEL, along with the two regulatory composites used as baselines. Section III details the preprocessing pipeline, graph construction,

Gabriel Luiz Espindola Pedro, Paulo Ricardo Branco da Silva, Arthur Cadore Matuella Barcella, João Vitor Rosa Negri, and Dimas Irion Alves are with Telecommunications Department of Aeronautics Institute of Technology, São José dos Campos, SP, Brazil, e-mails: {gabrieluizpe, dimasirion, negri, pbranco, arthurcadore, rmachado}@ita.br; Kai de Andrade Lima and Paulo Roberto Silva Marcos are with Agência Nacional de Telecomunicações (ANATEL), Brasília, Distrito Federal; e-mails: {kai.lima, paulorsm}@anatel.gov.br. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001, Anatel under Decentralized Execution Agreement (TED) No. 03/2023, Work Program: 24.722.2205.20ZD.0001, and São Paulo Research Foundation (FAPESP) under grant 2020/09838-0 (BIOS—Brazilian Institute of Data Science).

model architecture, and the multi-objective loss function that jointly drives the ranking optimization. Section IV evaluates the proposed model against the regulatory baselines through distribution and geographic analyses. Section V summarizes the main findings and outlines directions for future work.

II. DATASET AND BASELINE CREATION

The dataset comprises mobile service QoS and QoE indicators and infrastructure covariates provided by ANATEL, creating a matrix $\mathbf{X} \in \mathbb{R}^{N \times D}$, where N corresponds to the number of municipality–period entities and D corresponds to the number of features. The matrix $\mathbf{X}$ is indexed by two keys: *month*, which identifies the observation period, and a unique *municipality identifier*. Each row, therefore, corresponds to one municipality–month pair. Since not all selected features are available at monthly granularity, interpolation schemes are applied where necessary.

The eight indicators available related to mobile service are selected and normalized to $[0, 1]$: IND₁ (voice-call connection rate), IND₂ (voice-call drop rate), IND₃ (data-connection rate), IND₄ (download and upload speed compliance rate), IND₅ (data-connection latency), IND₆ (jitter), IND₇ (packet-loss rate), and IND₈ (service availability rate). In addition, INF_{4UP,{3G,4G,5G}} and INF_{4DL,{3G,4G,5G}} indicators are provided, which record the median upload and download speeds over each mobile data technology, respectively. For QoE, the ISG (General Satisfaction Index) was selected as a composite of QoE indicators for a given municipality–period tuple. Infrastructure is represented by the IBC (Brazilian Coverage Index).

To establish a baseline comparison, we adopt the two IQS composite scores¹. The IQS is officially published on a semi-annual basis, indexed by provider and municipality, with the primary objective of monitoring and evaluating telecommunications service quality in Brazil through a transparent and comparative indicator of provider performance [8], [9].

For our analysis, we derive monthly estimates aggregated at the municipality level by averaging across providers. The legacy composite $\text{QoS}_{\text{leg}} = \sum_{k=1}^K w_k^{\text{leg}} \text{IND}_k$, in use since November 2021, uses fixed weights $w^{\text{leg}} = (0.1, 0.1, 0.1, 0.2, 0.1, 0.1, 0.1, 0.2)$, assigning greater relevance to IND₄ and IND₈ [9]. The revised composite QoS_{rev} , replacing the legacy since June 2025, adopts a performance-dependent scheme: indicators are ranked from worst to best within each municipality–period and assigned the inverse-harmonic weight $\alpha_r = 1/(r H_K)$, where $H_K = \sum_{j=1}^K 1/j$, so that $\text{QoS}_{\text{rev}} = \sum_{r=1}^K \alpha_r \cdot \text{IND}_{\pi(r)} + B_{\text{cov}}$, with π consisting in a permutation that orders indicators from worst to best and B_{cov} consisting in a five-point bonus awarded when 4G or higher coverage reaches at least 95% in the urban area [8].

III. METHODOLOGY

The proposed ranking framework transforms raw municipality indicators into a scalar ranking score through pairwise

¹The IQS (Índice de Qualidade do Serviço) is referred to by ANATEL as the official QoS composite indicator.

Pareto dominance-guided learning on a graph-structured representation. The pipeline comprises four stages: preprocessing of raw indicators, detailed in Section III-A; graph construction in Section III-B; positional encoding in Section III-C; and a residual graph convolutional architecture trained with the multi-objective loss described in Section III-E.

A. Preprocessing

The selected features exhibit heavy-tailed distributions and marked heterogeneity in scale and support range—characteristics observed empirically across infrastructure and QoS indicators in the dataset. Algorithm 1 details the full preprocessing pipeline, which takes $\mathbf{X}$ as input and produces two output matrices: $\tilde{\mathbf{X}}$, used as input to the model, and the Pareto dominance proxy matrix $\mathbf{X}^{\mathcal{P}}$, used as a directional target during training. The pipeline first applies a Log1p transform, $x \mapsto \log(1 + x)$, to compress heavy tails. Next, Robust Scaling centers each feature by its median and divides by the interquartile range $Q_{0.95} - Q_{0.05}$, which is more resilient to outliers than the full range. Standard Scaling is then applied sequentially to yield $\tilde{\mathbf{X}}$ with zero mean and unit variance, ensuring that no single feature dominates gradient updates. In parallel, Min-Max scaling on $\mathbf{X}^{\log}$ produces $\mathbf{X}^{\mathcal{P}} \in [0, 1]^{N \times D}$, from which a scalar surrogate utility u_i is derived as the feature-wise mean; since all features are normalized to $[0, 1]$ and oriented so that higher values denote better performance, this average provides a consistent monotone direction for training in the absence of labeled rankings.

Algorithm 1 Preprocessing Pipeline

Input: $\mathbf{X} \in \mathbb{R}^{N \times D}$, raw indicator matrix

- 1: $\mathbf{X}^{\log} \leftarrow \log(1 + \mathbf{X})$ ▷ Log1p Transform
- 2: **for** each feature $j = 1, \dots, D$ **do**
- 3: $\tilde{x}_{:,j} \leftarrow \frac{\mathbf{x}_{:,j}^{\log} - \text{median}(\mathbf{x}_{:,j}^{\log})}{Q_{0.95}(\mathbf{x}_{:,j}^{\log}) - Q_{0.05}(\mathbf{x}_{:,j}^{\log})}$ ▷ Robust Scaling
- 4: **end for**
- 5: **for** each feature $j = 1, \dots, D$ **do**
- 6: $\tilde{\mathbf{X}}_{:,j} \leftarrow \frac{\tilde{x}_{:,j} - \text{mean}(\tilde{x}_{:,j})}{\text{std}(\tilde{x}_{:,j})}$ ▷ Standard Scaling
- 7: **end for**
- 8: **for** each feature $j = 1, \dots, D$ **do**
- 9: $\mathbf{X}_{:,j}^{\mathcal{P}} \leftarrow \frac{\mathbf{x}_{:,j}^{\log} - \min_i \mathbf{x}_{i,j}^{\log}}{\max_i \mathbf{x}_{i,j}^{\log} - \min_i \mathbf{x}_{i,j}^{\log}}$ ▷ Pareto Proxy
(Min-Max on $\mathbf{X}^{\log}$)
- 10: **end for**
- 11: **for** each entity $i = 1, \dots, N$ **do**
- 12: $u_i \leftarrow \frac{1}{D} \sum_j \mathbf{X}_{i,j}^{\mathcal{P}}$
- 13: **end for**

Ensure: $\tilde{\mathbf{X}} \in \mathbb{R}^{N \times D}$ (model input), $\mathbf{X}^{\mathcal{P}} \in [0, 1]^{N \times D}$, $u_i \in [0, 1]$ (ranking surrogate)

B. Graph Construction and Relational Structure

A directed k NN graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is constructed over the preprocessed feature space $\tilde{\mathbf{X}}$, where $\mathcal{V}$ is the set of nodes and $\mathcal{E}$ the set of directed edges. Each municipality–period is

represented by a node $v_i \in \mathcal{V}$, and edges encode feature-space proximity between entities.

The graph is constructed using k NN on $\tilde{\mathbf{X}}$ with $k = 15$ and Euclidean distance metric, yielding a neighbor matrix $\mathcal{I} \in \{1, \dots, N\}^{N \times k}$ where each row $\mathcal{I}_{i,:}$ lists the k nearest neighbors of node i . The resulting graph is directed: edges $(i \rightarrow j)$ are generated for each $j \in \mathcal{I}_{i,:}$.

C. Positional Encoding

Message-passing GNNs possess a permutation-equivariance property: relabeling nodes does not change the computation as long as edges and features are relabeled consistently [10]. While this invariance is generally desirable, it also means the network lacks the notion of absolute coordinates for its nodes. To inject this structural awareness, we compute positional encodings $p_i \in \mathbb{R}^{d_{pe}}$.

We employ Principal Component Analysis (PCA), which finds the orthogonal directions of maximum variance in $\tilde{\mathbf{X}}$ and projects onto the top components; each node thus receives a low-dimensional coordinate reflecting its position in the global data distribution.

While alternative methods—such as spectral embeddings based on the graph Laplacian or non-linear manifold techniques—can provide richer positional encodings, PCA is adopted here as a computationally efficient linear approximation suitable for the dataset scale.

The computed features p_i are standardized and concatenated with the original node data x_i , ensuring that the positional coordinates do not dominate as a result of disproportionate magnitudes.

D. Model Architecture

The core processing logic is built upon message-passing, generalizing the spectral graph convolution introduced in [11]. At each layer ℓ , a node representation $h_i^{(\ell)}$ is updated by aggregating features from its neighbors via

$$m_i^{(\ell)} = \text{AGG} \left(\left\{ \phi \left(h_j^{(\ell)} \right) : (j \rightarrow i) \in E \right\} \right), \quad (1)$$

with each subsequent layer computed as

$$h_i^{(\ell+1)} = \psi \left(h_i^{(\ell)}, m_i^{(\ell)} \right), \quad (2)$$

where ϕ is a learnable linear transformation, AGG is a permutation-invariant aggregator (mean pooling), and ψ combines the aggregated messages with the node's current state through a learnable affine transformation followed by a non-linear activation [12].

The proposed ranking framework uses Residual Graph Convolutional (ResGCN) blocks for ψ , incorporating skip-connections following the original residual learning paradigm [13], [14]. Given the node feature matrix $X \in \mathbb{R}^{n \times d_{in}}$ for a subgraph with n nodes and d_{in} input features, with the edge list stored in COO (Coordinate) format as source–target index arrays, each ResGCN layer computes

$$\begin{aligned} H &= XW_{\text{proj}}^\top + b_{\text{proj}}, \\ \bar{H}_i &= \frac{\sum_{(j \rightarrow i) \in E} H_j}{\deg_i + \epsilon}, \\ X' &= \text{GELU}(\bar{H} + \text{skip}(X)), \end{aligned} \quad (3)$$

where W_{proj} and b_{proj} are learnable weights and biases, $\deg_i$ is the in-degree of node i , and ϵ is a small stability constant. The skip path is the identity when dimensions match or a linear projection otherwise; this residual structure shortens gradient paths, improving trainability and mitigating oversmoothing—the convergence of node representations under repeated neighborhood aggregation [14], [15].

Due to the rapid computational growth associated with neighborhood scaling, bounded graph subsampling via breadth-first search with neighbor downsampling is employed. At each hop, neighbor lists are randomly subsampled to a fixed limit of 5, and a hard cap on the total number of nodes is enforced. This enables memory-efficient training by transferring only small induced subgraphs to the processing device at each step, at the cost of reduced long-range message propagation per step, which stochastic sampling across iterations mitigates.

All activations use the Gaussian Error Linear Unit (GELU), achieving smooth, continuously differentiable non-linearity. The GELU activation function is monotone increasing for positive and small negative values; coupled with the model components, this property satisfies monotonicity constraints: positive variations in individual QoS metrics consistently translate to non-decreasing scores.

The model functions as a multi-task architecture with two output heads from the final representations: (1) a *ranking head* that maps the node embedding to a scalar score $\hat{r}_i \in [0, 1]$ via an MLP with sigmoid output; and (2) a *reconstruction head* that maps the embedding back to the original feature space. The reconstruction head acts as an autoencoder regularizer, discouraging degenerate solutions where ranking constraints are satisfied but the intermediate representation discards input information. Each of these objectives is formalized through a dedicated loss term, as described in the following subsection.

E. Training Objectives and Loss Terms

The model uses a weighted sum of four loss terms to jointly optimize complementary perspectives, yielding a scalarized solution more balanced and robust than one-dimensional optimization

$$\mathcal{L} = \lambda_{\text{comp}} \mathcal{L}_{\text{comp}} + \lambda_{\text{recon}} \mathcal{L}_{\text{recon}} + \lambda_{\text{rank}} \mathcal{L}_{\text{rank}} + \lambda_{\text{geo}} \mathcal{L}_{\text{geo}}, \quad (4)$$

The Compass Loss $\mathcal{L}_{\text{comp}}$ aligns model scores with the proxy u_i , providing a monotone directional signal during training

$$\mathcal{L}_{\text{comp}} = \frac{1}{|B|} \sum_{i \in B} (s_i - u_i)^2. \quad (5)$$

where i indexes each sample in the current mini-batch B .

The Reconstruction Loss $\mathcal{L}_{\text{recon}}$ evaluates the decoder’s ability to recover the original standardized features, enforcing information preservation

$$\mathcal{L}_{\text{recon}} = \frac{1}{|B|} \sum_{i \in B} \|\hat{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|_2^2. \quad (6)$$

The Pairwise Ranking Loss $\mathcal{L}_{\text{rank}}$ establishes ordering constraints via margin-based pairwise comparisons: the loss penalizes any pair whose predicted score gap falls below a minimum margin m relative to the ordering induced by the proxy. This follows the classical RankNet formulation [16] and its neural extensions [17]. For two entities A and B with target label $y = \text{sign}(u_A - u_B)$, where u_A and u_B are the proxy utility values of municipalities A and B , respectively, we have that

$$\mathcal{L}_{\text{rank}} = \frac{1}{|P|} \sum_{(A,B) \in P} \max(0, m - y(s_A - s_B)), \quad (7)$$

where the margin $m = 0.1$ serves a dual purpose: it demands a minimum score separation between pairs ordered by the proxy, and it suppresses the contribution of pairs whose proxy difference $|u_A - u_B|$ is small—and therefore potentially unreliable—by allowing the loss to reach zero before perfect separation is achieved. Pairs that violate the margin constraint are penalized linearly, while correctly ordered pairs with sufficient separation contribute zero loss.

The Geometric Smoothness Loss $\mathcal{L}_{\text{geo}}$ encourages score consistency among neighboring nodes by penalizing large score differences across edges, following the discrete Dirichlet energy—the sum of squared signal differences along all edges—on the graph

$$\mathcal{L}_{\text{geo}} = \frac{1}{|E_{\text{sub}}|} \sum_{(i,j) \in E_{\text{sub}}} (s_i - s_j)^2, \quad (8)$$

where E_{sub} is the edge set of the current subgraph sample. Since graph edges connect municipalities that are similar in feature space, this term enforces that nearby entities receive comparable scores. A conservative weight prevents oversmoothing, preserving meaningful score variation across the distribution.

IV. RESULTS AND DISCUSSION

We evaluate the proposed ranking framework against the regulatory baselines using distribution and geographic analyses across 82,462 municipality–period observations. The legacy composite QoS_{leg} concentrates scores in a narrow band ($\sigma = 0.054$, Interquartile Range (IQR) = 0.028), and the revised QoS_{rev} offers only marginal improvement ($\sigma = 0.058$, IQR = 0.028). In contrast, the proposed score $\hat{r}$ achieves $\sigma = 0.098$ and IQR = 0.137—a 1.8 times and nearly five times expansion relative to QoS_{leg} , respectively—enabling finer municipality rank differentiation with distinct performance profiles.

Fig. 1 presents overlaid histograms of the three scoring approaches, enabling direct visual comparison of how each method distributes municipality scores. The QoS_{leg} distribution is comparatively more clustered: its mass is almost




Fig. 1: Comparison between score and baseline distributions.

entirely concentrated between 0.83 and 0.87, forming a sharp, narrow peak that reflects the inability of the legacy composite to spread municipalities across the score domain. The QoS_{rev} distribution is slightly wider but displays a pronounced asymmetry, with a tall peak near 0.59 and a long left tail extending toward 0.05. In contrast, the proposed model’s histogram is visibly broader and flatter, with frequencies distributed more uniformly across the range from approximately 0.22 to 0.80. This shape indicates that the multi-objective optimization yields a scoring function that more effectively utilizes the support domain, producing a more differentiated range of performance levels.

As external validation, the proposed score achieves a Pearson correlation of $\rho = 0.40$ with the Municipal Human Development Index (MHDI)—independent of the intra-set features—exceeding the baselines by 0.07 and 0.05, indicating that the model captures socioeconomic structure not encoded in the linear composites.

To spatially validate the framework, we selected December 2022 for evaluation, as it corresponds to the period with the fewest missing values across municipalities. Fig. 2 shows the comparative geographic mapping for each metric across Brazilian municipalities. While the proposed model score exhibits continuous performance transitions that highlight distinct geographic regions—consistent with the distributional patterns observed in the histogram analysis—the regulatory indicators collapse the vast majority of municipalities into either the highest score band or a narrow intermediate range.

The geographic contrast is particularly evident when comparing macro-regions. In the QoS_{leg} map, municipalities with substantially different performance profiles are assigned nearly identical composite values, creating plateaus that flatten the distribution—masking known infrastructure and socioeconomic disparities between these regions. According to the MHDI compiled by the United Nations Development Programme, all states in the North and Northeast present MHDI values below the national average, in stark contrast to the South and Southeast, where the majority of municipalities fall within the higher development tiers [18]. The QoS_{rev} map, despite its inverse-harmonic weighting scheme, still exhibits limited score dispersion, with most of the territory concentrated in a narrow range—even though broad regional contrasts, such as the separation between the North and the remaining regions, remain discernible. The proposed framework




Fig. 2: Geographic distribution of analysed variables for December 2022.

amplifies this variation, extending discriminability across the full score range.

The proposed model’s score map reveals a clear North–South gradient consistent with the known human development landscape across Brazilian macro-regions [18]. Municipalities in the South, Southeast, and parts of the Center-West display higher scores, forming distinguishable clusters. In comparison, the North and parts of the Northeast exhibit lower scores aligned with their lower MHDl scores [18]. Even within macro-regions, the model captures intra-regional heterogeneity: isolated high-performing municipalities (typically state capitals or large urban centers) are clearly visible against their surrounding rural areas, a distinction that the baselines’ compressed ranges cannot render.

V. CONCLUSION

While the analyzed ANATEL composite indices exhibit score clustering that limits spatial discrimination, this paper introduces a self-supervised, Pareto-guided ranking framework that significantly expands score discriminability across the distribution of municipalities, surpassing legacy and revised baselines through joint optimization of competing quality objectives.

By combining pairwise ranking, feature reconstruction, geometric smoothness, and a directional penalty, the model’s optimization covers distinct, complementary dimensions, yielding a more balanced solution than that provided by one-dimensional scalar optimization. The quantitative analysis demonstrates that the proposed score achieves an interquartile range of 0.137, representing a nearly five times expansion compared to both baselines—IQR = 0.028 for QoS_{leg} and IQR = 0.028 for QoS_{rev} —and a standard deviation 1.8 times larger than the QoS_{leg} . External validation against the MHDl shows that the proposed score correlates with MHDl at $\rho = 0.40$, exceeding the correlation achieved by the legacy and

revised baselines by 0.07 and 0.05, respectively, indicating that the learned ranking captures socioeconomic structure beyond what the linear composites encode.

Future work can benchmark the proposed approach against established learning-to-rank methods. Strengthening the external validation with additional covariates outside the current dataset and seemingly independent of telecommunications metrics is also a possible direction of research. Finally, exploring alternative positional encoding methods, such as Spectral Embedding for non-linear manifold coordinates, and testing other activation functions to assess infrastructural stability are also possibilities.

REFERENCES

- [1] M. Liu, H. Xu, Q. Z. Sheng, and Z. Wang, “Qosgnn: Boosting qos prediction performance with graph neural networks,” *IEEE Transactions on Services Computing*, vol. 17, no. 2, pp. 645–658, 2024.
- [2] W. Wang, W. Ma, and K. Yan, “Trust-aware privacy-preserving qos prediction with graph neural collaborative filtering for internet of things services,” *Complex & Intelligent Systems*, vol. 11, no. 4, p. 191, 2025.
- [3] X. Zhang and L. Li, “Assessing the quality of experience in wireless networks for multimedia applications: A comprehensive analysis utilizing deep learning-based techniques,” *Heliyon*, vol. 10, no. 9, p. e30351, 2024.
- [4] Y. Li, Z. Lyu, Y. Zhang, H. Zhang, T. Peng, H. Xiong, S. Wang, L. Kong, G. Chen, and D. Yin, “S³prank: Toward satisfaction-oriented learning to rank with semi-supervised pre-training,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 38, no. 1, pp. 559–572, 2026.
- [5] A. Suresh and K. Deb, “Machine learning-based prediction of new pareto-optimal solutions from pseudo-weights,” *IEEE Transactions on Evolutionary Computation*, vol. 28, no. 5, pp. 1351–1365, 2024.
- [6] A. Kesireddy and F. A. Medrano, “Elite multi-criteria decision making—pareto front optimization in multi-objective optimization,” *Algorithms*, vol. 17, no. 5, 2024.
- [7] E. Haxhiraj, D. Shahu, and E. Agastra, “Pareto front transformation in the decision-making process for spectral and energy efficiency trade-off in massive mimo systems,” *Sensors*, vol. 25, no. 5, 2025.
- [8] Agência Nacional de Telecomunicações, “Resolução interna anatel n° 444,” 2025.
- [9] —, “Resolução interna anatel n° 71,” 2021, revogada.
- [10] Y. Dai, L. Lyu, N. Cheng, M. Sheng, J. Liu, X. Wang, S. Cui, L. Cai, and X. Shen, “A survey of graph-based resource management in wireless networks—part ii: Learning approaches,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 4, pp. 2101–2122, 2025.
- [11] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *The 5th International Conference on Learning Representations (ICLR)*, Toulon, France, Apr. 2017.
- [12] A. Amara, M. A. Hadj Taieb, and M. Ben Aouicha, “A multi-view gnn-based network representation learning framework for recommendation systems,” *Neurocomputing*, vol. 619, p. 129001, 2025.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [14] M. Scholkemper, X. Wu, A. Jadbabaie, and M. T. Schaub, “Residual connections and normalization can provably prevent oversmoothing in GNNs,” in *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.
- [15] S. Akansha, “Over-squashing in graph neural networks: A comprehensive survey,” *Neurocomputing*, vol. 642, p. 130389, 2025.
- [16] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender, “Learning to rank using gradient descent,” in *Proceedings of the 22nd International Conference on Machine Learning*, ser. ICML ’05. New York, NY, USA: Association for Computing Machinery, 2005, p. 89–96.
- [17] M. Köppel, A. Segner, M. Wagener, L. Pensel, A. Karwath, and S. Kramer, “Pairwise learning to rank by neural networks revisited: Reconstruction, theoretical analysis and practical performance,” *Machine Learning*, vol. 114, no. 4, p. 112, 2025.
- [18] United Nations Development Programme, “Municipal human development index (MHDl) panel – Brazil,” UNDP Brazil Human Development Unit, 2023, accessed: 2025.