A Comparative Study of Gravity-Based Centrality Models for Ranking Autonomous Systems in Directed Acyclic Graphs

JACCO TE POEL, University of Twente, The Netherlands

This paper explores the use of gravity-based centrality measures to rank nodes in Directed Acyclic Graphs (DAGs) of Autonomous System (AS) networks. These measures are inspired by the classical gravity model, where influence is determined by node size and distance. We compare twelve variants of gravity-based centrality, each with distinct weighting and distance formulations, against customer cone size, a widely used AS ranking metric. Using historical AS DAG snapshots from 1998 to 2025, we evaluate each method in terms of ranking monotonicity, correlation, and stability. The results show that several gravity-based models, particularly DKEGM, Entropy-Based, and Multi-Characteristics Gravity, consistently outperform customer cone size in ranking resolution and stability. After removing leaf nodes, these models achieve near-perfect monotonicity ($\approx$ 0.9999) and strong agreement among themselves, offering a more nuanced view of AS importance in directed network topologies.

Additional Key Words and Phrases: Autonomous System Network, Autonomous System Ranking, Gravitational ranking

1 INTRODUCTION

The internet is made up of thousands of smaller networks, called Autonomous Systems (ASes). An AS is a collection of IP networks and routers under the control of a single organization that presents a unified routing policy to the rest of the internet [5]. Each AS is responsible for forwarding packets correctly within its network and to neighboring ASes.

Some examples of ASes are Internet Service Providers (ISPs), universities, large enterprises, and government networks[3]. While each AS operates independently, it participates in global routing by exchanging information with other ASes, using the Border Gateway Protocol (BGP).

When multiple of these ASes are connected, they form an Autonomous System Network (ASN). This network is created when multiple of these ASes exchange routing information using the Border Gateway Protocol (BGP), enabling end-to-end data transfer across different ASNs. The ASN shows relationships between ASes and defines how data is routed from one AS to another. The structure of the ASN changes dynamically based on routing policies, peering agreements.

The structure of an ASN can be represented as a Directed Acyclic Graph (DAG), where each node is an AS and each edge is a valid route between them. Since routing loops are avoided in BGP, the resulting topology forms a DAG[2]. This structure helps analyze routing behavior, identify influential ASes, and apply graph-based algorithms.

Edges in the DAG represent different types of inter-AS relationships:

- **Provider-to-Customer (p2c)**: A provider allows its customer to send and receive traffic through its network.
- **Customer-to-Provider (c2p)**: A customer connects through a provider to reach the internet.
- **Peer-to-Peer (p2p)**: Two ASes exchange traffic between their customers without paying each other.

Although the DAG edges represent route advertisements, actual data packets may flow in either direction, depending on the policies of intermediate ASes. For example, an AS at the bottom of the DAG may reach another AS on a different branch if providers along the way permit the traffic. Thus, the DAG is a logical abstraction used to show structure and policy relationships, not a strict enforcement of data flow direction[4].

1.1 Contributions

This paper makes the following key contributions:

- We provide a comparative analysis of twelve gravity-based centrality models applied to Directed Acyclic Graphs (DAGs) representing Autonomous System (AS) networks.
- We propose an evaluation framework that uses monotonicity and rank correlation to assess the effectiveness of these models in separating and ranking AS nodes.
- We demonstrate that certain gravity-based models (e.g., DKEGM, Multi-Characteristics, and Entropy-based) outperform the widely used customer-cone size metric in both ranking resolution and structural insight.
- We identify structural differences between gravity-based models and show how they reveal ASes with strategic importance not captured by traditional metrics.

To the best of our knowledge, this is the first work that systematically compares multiple gravity-based centrality models on large-scale DAGs of AS networks.

1.2 Motivation

1.2.1 Why rank ASes? Ranking ASes is important for businesses that operate their own networks. In particular, it plays an important role when establishing peer-to-peer (p2p) connections, which are usually created to improve local or regional routing efficiency. By identifying and ranking ASes based on their influence or centrality in the network, a business can make more informed decisions about which ASes to peer with, ultimately maximizing connectivity, performance, and resilience across the internet.

1.2.2 Role of centrality in understanding network influence and dynamics. Centrality is used to measure how important or impactful a node is in a network. For ASNs, centrality is used to find important ASes for connectivity in a network. These ASes usually have more connections or appear on many paths between other ASes, making them important for forwarding traffic. Ranking ASes based on centrality gives insight into the structure of the network and shows how resilient or efficient it is. It can also be used to find bottlenecks, improve routing, or decide where to create new peerings.

1.2.3 Limitations of customer-cone and other traditional approaches. There are several traditional ranking methods used to evaluate AS importance, such as Degree Centrality, Eigenvector Centrality, Alpha Centrality, customer-cone size, and Betweenness Centrality

TScIT 43, July 4, 2025, Enschede, The Netherlands

[16]. While these methods are useful, each has its own limitations when applied to ASNs.

- **Degree Centrality**: Measures how many direct connections an AS has. It does not consider the quality or importance of those connections [13].
- **Eigenvector Centrality**: Accounts for the influence of neighboring nodes, but assumes symmetric relationships, which is often not the case in ASNs [6].
- **Alpha Centrality**: Similar to eigenvector centrality but includes external influence. However, it requires tuning a parameter and lacks clear interpretation in some network types [1].
- **Customer-cone size**: Counts the number of ASes that can be reached through provider-to-customer paths. It ignores peering and c2p links, which might result in different ASes receiving similar rank indexes, potentially reducing ranking resolution [12].
- **Betweenness Centrality**: Measures how often an AS appears on the shortest paths between others. It's computationally expensive and can overemphasize nodes that lie between sparse regions of the graph [13].

1.3 Problem Statement

1.3.1 Problem Description. Traditional methods for ranking ASes often rely on standard centrality measures, such as degree or betweenness, which do not fully show the structure of ASNs. These metrics either oversimplify connectivity or are too computationally expensive for large-scale graphs. Gravity-based centrality is a more flexible approach that uses node properties (like size or distance) to model influence. However, there are multiple ways to define and compute gravity in network graphs, and it is unclear which variation works best for DAGs like ASNs.

1.3.2 Research Questions.

- How do different gravity-based centrality models, specifically the basic gravity model, improved gravity, DKEGM, Laplacian gravity, generalized gravity, local gravity, k-shell-based gravity, multi-characteristics gravity, and the entropy-based gravity model, perform on ASNs in terms of monotonicity and stability?
- Can gravity-based centrality offer a more scalable and discriminative alternative to traditional AS ranking methods like degree, betweenness, and customer-cone size?
- Can gravity-based centrality enhance the customer-cone size metric, which is currently the standard for ranking ASes, to more accurately reflect AS influence in DAGs?

2 STATE OF THE ART

Ranking ASes in ASNs is not a new concept. Many methods have been used, like degree, eigenvector, betweenness, and customer-cone size. Customer-cone is still the most common [12], but it only counts downstream ASes in p2c links, ignoring peering. This can give a skewed view in flat or hybrid topologies.

Standard graph centralities like degree are fast but only look at direct neighbors. Eigenvector and alpha centrality weight important neighbors more, but they assume symmetric links, which is not true for ASNs [1]. Betweenness looks at shortest paths but becomes slow on large networks [16].

To fix these issues, gravity-based models have been introduced. They combine a node's "mass" with how far away it is from others. Most models follow the idea of Newton's law. The basic gravity model uses degree or k-shell as mass [10]. The improved gravity model (IGM) [17] limits influence with a radius to reduce noise. The generalized gravity model (GGM) [7] adds local clustering to the mass to reflect neighborhood structure. Local gravity (LGM) [9] focuses on short-range influence using degree and distance.

Other versions go further. K-shell based gravity (KSGM) [18] boosts mass using exponential weights on the k-shell score. The multi-characteristics gravity model (MCGM) [8] mixes normalized values of degree, k-shell, and eigenvector into one term. The entropy-based gravity model (SEGM) [11] uses local entropy to adjust mass, making it sensitive to uneven connectivity. Finally, DKEGM [15] combines degree, k-shell, and eigenvector into a score.

Most of these models have only been tested on undirected or spatial graphs. But ASNs are directed, hierarchical, and often tree-like. So far, few papers look at how gravity works on DAGs. This research fills that gap by comparing multiple gravity models on AS DAGs and testing how they hold up against traditional metrics.

3 BACKGROUND

This section describes the DAG structure and the centrality metrics used in our analysis: degree centrality, k-shell decomposition, and eigenvector centrality. Specific changes have been made to address common issues when applying these methods to DAGs.

3.1 DAGs

A DAG is a directed graph that contains no cycles. Formally, let $G = (V, E)$ be a graph with vertex set V and directed edge set $E \subseteq V \times V$. The graph G is a DAG if and only if it contains no directed cycles, i.e., there does not exist a sequence of vertices $v_1, v_2, \ldots, v_k \in V$ with $k \geq 2$ such that:

$$(v_i, v_{i+1}) \in E \quad \text{for all } i = 1, 2, \ldots, k-1 \quad \text{and} \quad (v_k, v_1) \in E.$$

3.2 Degree Centrality

In a directed graph, each node $v \in V$ has an in-degree k_v^{in} and an out-degree k_v^{out}, defined as:

$$k_v^{\text{in}} = |\{u \in V : (u, v) \in E\}| \tag{1}$$

$$k_v^{\text{out}} = |\{w \in V : (v, w) \in E\}| \tag{2}$$

To avoid problems with nodes that have either zero in-degree or out-degree, which can result in division-by-zero errors in some models, we use the total degree:

$$k_v = k_v^{\text{in}} + k_v^{\text{out}} \tag{3}$$

This combined metric ensures that each node has a non-zero degree value in most cases, allowing consistent and robust processing.

3.3 K-shell Decomposition

K-shell decomposition assigns nodes to layers (or shells) based on their degree [14]. In a directed setting, one can define in-shells or out-shells using k^{in} or k^{out}, but this is problematic in DAGs where many nodes naturally have zero in- or out-degree.

To avoid such instability and division-by-zero issues, we again use the total degree k_v as defined above. The algorithm recursively removes all nodes with $k(\text{degree}) \leq ks$ until no such nodes remain, and assigns those nodes to shell ks.

3.4 Eigenvector Centrality

Eigenvector centrality [6] is based on the idea that connections to important nodes contribute more to a node's importance. Formally, it is defined as the solution to:

$$Ax = \lambda x \tag{4}$$

where A is the adjacency matrix of the graph, λ is the largest eigenvalue, and x is the corresponding eigenvector.

In the case of a DAG, using the directed adjacency matrix A leads to $x = 0$, as there are no cycles. To address this, we transform the DAG into an undirected graph by treating all directed edges as undirected, and construct a new adjacency matrix A^* for this modified structure:

$$A^*_{ij} = \begin{cases} 1 & \text{if } (i,j) \in E \text{ or } (j,i) \in E \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

We then compute eigenvector centrality on this undirected matrix, which yields meaningful, non-zero scores for all nodes. Here, x_v denotes the centrality score of node v, with higher values indicating greater importance within the network.

4 METHODOLOGY

Symbol	Definition
k_i	Degree of node i
$k_{\max}$	Maximum degree in the network
k_{mid}	Median degree in the network
ks_i	K-shell value of node i
$ks_{\max}$	Maximum k-shell value in the network
$ks_{\min}$	Minimum k-shell value in the network
ks_{mid}	Median k-shell value in the network
x_i	Eigenvector centrality of node i
$x_{\max}$	Maximum eigenvector centrality in the network
x_{mid}	Median eigenvector centrality in the network
A_i	Set of neighbors of node i
n_i	The number of edges between neighbors of node i.
R	Truncation radius (maximum distance considered)
d_{ij}	Shortest path distance between nodes i and j
λ	Eigenvalue of Laplacian matrix

Table 1. Symbol definitions used in gravity-based centrality models

4.1 Model Implementation

4.1.1 Gravity Model. The gravity model [10] measures influence by combining how "massive" a node is (its k-shell index) with how close it is to other nodes. Based on Newton's law, two nodes influence each other more if they're both strong and nearby. The gravity of node i is calculated as:

$$GM(i) = \sum_{d(i,j),\ i \neq j} \frac{ks_i \cdot ks_j}{d_{ij}^2} \tag{6}$$

This means that two high-k nodes just a few hops apart contribute more to each other's score. The effect drops off quickly with distance.

The extended version also includes influence from the gravity scores of direct neighbors:

$$GM_+(i) = \sum_{n \in A_i} GM(n) \tag{7}$$

This adds local structure to the ranking and helps reduce high scores from nodes that are structurally unimportant.

4.1.2 Improved Gravity Model. To further improve the model, a truncation radius is used [17]:

$$IGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{ks_i \cdot ks_j}{d_{ij}^2} \tag{8}$$

An optimal truncation radius R^* can be estimated using:

$$R^* \approx \frac{1}{2}\langle d \rangle \tag{9}$$

where $\langle d \rangle$ is the average shortest path length in the network.

Similar to the gravity model, there is also an extended version of the Improved Gravity Model [17]:

$$IGM_+(i) = \sum_{n \in A_i} IGM(n) \tag{10}$$

The idea is the same as the normal gravity model, but now we control how far the influence spreads. Instead of summing over all nodes, we cut it off at a fixed range R, so only closer ASes count. This makes the model more scalable and keeps far-away weak links from inflating the score. The extended version again includes influence from nearby nodes.

4.1.3 Local Gravity Model. The local gravity model [9] is defined as follows:

$$LGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{k_i \cdot k_j}{d_{ij}^2} \tag{11}$$

This model introduces a truncation radius R, just like the improved gravity model, but instead of k-shell, it uses the degree of nodes to represent their mass. This makes the model faster to compute and more sensitive to local structure. The idea is that influence should mostly come from close neighbors, and nodes with more links contribute more mass. By limiting interactions to nearby ASes, it reduces noise from distant low-impact nodes.

4.1.4 Generalized Gravity Model. The generalized gravity model [7] builds on the gravity idea but introduces an additional refinement. Instead of using raw degree or k-shell as mass, it defines a node's "spreading ability" Sp_i, which is based on how well it can spread info to others not just how many neighbors it has, but also how tightly those neighbors are clustered.

$$GGM(i) = \sum_{d(i,j),\ i \neq j} \frac{Sp_i \cdot Sp_j}{d_{ij}^2} \tag{12}$$

Where Sp_i is:

$$Sp_i = e^{-2C_i} \times k_i \tag{13}$$

Where C_i is the clustering coefficient:

$$C_i = \frac{2n_i}{k_i(k_i - 1)} \tag{14}$$

The underlying assumption is that nodes with a high number of connections and whose neighbors are not just connected to each other is likely to spread influence further. If its neighbors are very connected among themselves, information gets trapped locally so we reduce that node's weight.

4.1.5 K-shell Based Gravity Model. The k-shell based gravity model [18] builds on the idea that nodes deeper in the core of the network (with higher k-shell values) are more stable and densely connected. To reflect this, the model adjusts the gravitational interaction between nodes using a correction factor c_{ij}, which scales the influence based on how central both nodes are.

$$KSGM(i) = \sum_{d(i,j),\ i \neq j} c_{ij} \frac{k_i \cdot k_j}{d_{ij}^2} \tag{15}$$

$$c_{ij} = e^{\frac{ks_i ks_j}{ks_{max} - ks_{min}}} \tag{16}$$

When two nodes are both in the core, c_{ij} boosts the gravity. If they are on the edge of the core, the boost is smaller. So this model gives more influence to nodes that are better positioned in the network structure.

4.1.6 DKEGM Gravity Model.

The DKEGM model [15] attempt to address the limitations of traditional measures (like only using degree or k-shell in the mass) in distinguishing between nodes. A lot of ASes may have the same values in those simple metrics. So, this model combines three things: degree, eigenvector centrality, and an improved k-shell score into one index called DKE.

This index gives each node a more distinct score, which helps avoid ties. The DKE score is used as the mass in the gravity equation:

$$DKE(i) = k_i + ks_i^* + x_i \tag{17}$$

$$ks_i^* = ks_i + \frac{R(i)}{\max_k S(k) + 1} \tag{18}$$

where the overall count of stages is $S(k)$ during the k-shell decomposition for a specific k-degree iteration, and $R(i)$ is in which stage i is eliminated.

Using this more detailed score, the model applies gravity with a truncation radius:

$$DKEGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{DKE(i) \cdot DKE(j)}{d_{ij}^2} \tag{19}$$

By doing this, the model assigns higher importance to nodes that score highly across different centrality measures and are also close to each other. It's especially useful when you want to better distinguish key nodes in large networks where many nodes would otherwise receive similar scores.

4.1.7 Laplacian Gravity Model.

The Laplacian gravity model [19] swaps out node degree for something more refined: Laplacian centrality. This centrality does not just count how many links a node has but also looks at how its removal changes the entire structure, giving a better picture of its overall importance.

The Laplacian matrix can be defined as follows:

$$L = D - A \tag{20}$$

where D is the degree matrix and A is the adjacency matrix.

The energy of the full graph and the graph with node i removed are:

$$E(G) = \sum_{\lambda \in L} \lambda^2 \quad E(G_{-i}) \quad = \sum_{\lambda \in L_{-i}} \lambda^2 \tag{21}$$

The centrality is the difference in energy:

$$LC(i) = E(G) - E(G_{-i}) \tag{22}$$

Finally, this is plugged into a gravity model:

$$LPGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{LC(i) \cdot LC(j)}{d_{ij}^2} \tag{23}$$

This approach uses spectral information from the graph, making it more sensitive to both local and global structure. It captures not just how connected a node is, but how important it is to the network's overall structure.

4.1.8 Multi-characteristics Gravity Model.

The multi-characteristics gravity model [8] builds on the idea that a single metric is not enough to capture a node's true importance. Instead of using just degree or k-shell, this model combines three: degree centrality, k-shell index, and eigenvector centrality each measuring something different about a node's position and influence in the network.

Since these metrics can differ in scale, they're normalized and combined with a correction factor α, which balances the smaller k-shell range compared to the others. The model then applies the gravity formula using this combined metric.

$$w_i = \frac{k_i}{k_{\max}} + \frac{\alpha ks_i}{ks_{\max}} + \frac{x_i}{x_{\max}} \tag{24}$$

$$MCGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{w_i w_j}{d_{ij}^2} \tag{25}$$

$$\alpha = \frac{\max\{\frac{k_{mid}}{k_{max}}, \frac{x_{mid}}{x_{max}}\}}{\frac{ks_{mid}}{ks_{max}}} \tag{26}$$

This way, the model captures multiple layers of a node's network position from local connections to global influence and reduces bias from using only one metric.

4.1.9 Entropy-based Gravity Model.

The entropy-based gravity model [11] improves on previous gravity approaches by using information entropy to better capture how influence is spread across a node's neighborhood.

$$SEGM(i) = \sum_{d(i,j) \leq R,\ i \neq j} \frac{SE(i) \cdot SE(j)}{d_{ij}^2} \tag{27}$$

$$SE(i) = e^{E(i)} k_i \tag{28}$$

$$E(i) = -\sum_{j \in A_i} I(j) \ln I(j) \tag{29}$$

$$I(i) = \frac{k_i}{\sum_{j \in A_i} k_j} \tag{30}$$

The idea behind this model is that influence is not just about having a lot of neighbors it also depends on how evenly distributed the neighbor degrees are. If a node's neighbors have similar degree values, the uncertainty (entropy) is higher. This means that there is a stronger ability to spread information. This model combines that entropy with gravity, giving more importance to nodes that are both well-connected and structurally diverse.

4.2 Dataset Usage

The dataset we use comes from CAIDA's AS Relationships dataset [2]. It represents the ASN structure based on BGP routing data. We treat the network as a DAG, where each node is an AS and each directed edge is a p2c link.

We collected all CAIDA AS relationship snapshots from January 1, 1998 through May 1, 2025. This range covers many years of Internet growth, ensuring each model is evaluated on a wide variety of network sizes and structures.

The May 2025 snapshot (the latest in our set) contains 401,699 nodes and 160,292 directed edges. The average degree is around 0.3990, meaning each AS has on average less than one connection in this directed graph (reflecting the many leaf ASes). The graph's diameter is 25 (the longest shortest path between two ASes), and the density is very low (close to 0), as expected for a large-scale Internet network.

We focus on the largest connected component to make sure the network is fully reachable. This structure is important when computing centrality scores, especially gravity-based models that depend on distance and connectivity across the graph.

4.3 Evaluation Metrics

To compare the different gravity-based centrality models, we employ several evaluation methods. These assess not only how distinct the rankings are within each model but also how consistent they are with one another and with known indicators of influence in Autonomous System (AS) networks, such as customer-cone size.

4.3.1 Ranking Resolution (Monotonicity). Ranking resolution, also referred to as monotonicity, measures how well a centrality model distinguishes between nodes. A model with high resolution will produce fewer tied ranks and provide clearer separation across the ranking scale.

We compute monotonicity using the following formula:

$$M(X) = [1 - \frac{\sum_{c \in V} N_c(N_c - 1)}{N(N-1)}] \tag{31}$$

where N_c is the number of nodes with the same rank and N is the size of the entire network.

4.3.2 Rank Correlation. To evaluate the correlation between rankings produced by the different models, we compute two standard rank correlation coefficients: Spearman's ρ and Kendall's τ.

Spearman's rank correlation coefficient is defined as:

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)} \tag{32}$$

where d_i is the difference in ranks of node i between the two models, and n is the total number of nodes.

Kendall's Tau is computed as:

$$\tau = \frac{n_c - n_d}{\frac{1}{2} n(n-1)} \tag{33}$$

where n_c is the number of concordant node pairs and n_d is the number of discordant pairs. Both ρ and τ range from −1 (complete disagreement) to 1 (perfect agreement), with higher values indicating stronger alignment between the compared centrality models.

4.3.3 Ranking Comparison. To better understand how the different models rank influential ASes, we generate a table for each centrality model alongside the customer-cone ranking. These tables display the top-n ASes according to each approach, allowing for a direct comparison of their outputs.

With this, we can observe how rankings differ across models for example, if ASes with lower rank in customer-cone score higher in certain gravity-based models. Such differences may indicate that a model captures different aspects of network influence beyond downstream reach, such as if there is a well-connected node is in its neighborhood.

5 RESULTS

In some of the results, leaf nodes are removed before calculating correlation and monotonicity to improve the quality of the evaluation. This is done because leaf nodes usually receive a score of zero in both the customer-cone and gravity-based models. Including them can give a higher correlation and a lower monotonicity.

In Figure 1, x_4, x_5, and x_6 are leaf nodes. Under the customer-cone model, these nodes have no customers below them, so their

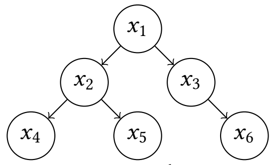

Fig. 1. Example DAG

cone size is always zero. In the gravity-based model, the score is also zero. The gravity score between two nodes i and j is calculated as:

$$\text{Gravity}(i,j) = \frac{w_i \cdot w_j}{\text{dist}(i,j)^2}$$

where w_i and w_j are node weights and $\text{dist}(i,j)$ is the path length between them. For a leaf node, there are no reachable neighbors, so no valid paths exist, and the total gravity score becomes:

$$\sum_{j \neq i} \frac{w_i \cdot w_j}{\text{dist}(i,j)^2} = 0$$

To account for this, we also present results with leaf nodes excluded.

5.1 Monotonicity

Figure 2 and 3 show the monotonicity scores for all the centrality models. Monotonicity measures how consistent the ranking is with the graph's topological structure. A high score means the model is better at producing a ranking that follows the DAG's flow.

With leafs: As seen in Figure 2, all models score below 0.1 when we include leaf nodes. This is expected, since leaves do not have any outgoing connections, so their centrality values end up being 0 or close to each other. This flattens the overall ranking and makes it hard to distinguish between nodes. In fact, models like Extended Gravity, Generalized Gravity, and Laplacian Gravity have monotonicity values near 0 in this case, because they give nearly identical (zero) scores to most of those leaf ASes.

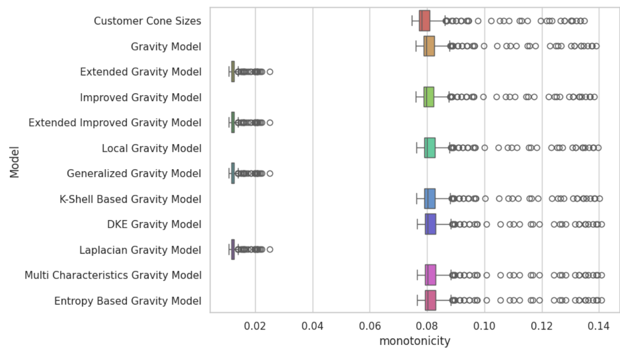

Fig. 2. Monotonicity scores across all snapshots **with leaf nodes included**. Most models perform poorly due to tied zero scores from leaves, which flatten the rankings.

Without leafs: Removing the leaves changes the picture completely (Figure 3). DKE, Multi-Characteristics, and Entropy-Based Gravity models now score nearly 1.0 (essentially perfect). K-Shell Gravity is above 0.97. Gravity, Local Gravity, and Improved Gravity all reach around 0.90-0.95. Even the Customer-Cone metric improves to about 0.75 without leaves (higher than before but still below the top gravity models). In contrast, Extended Gravity, Generalized Gravity, and Laplacian Gravity remain far behind, only around 0.35, meaning they still produce many ties and inconsistent ordering.

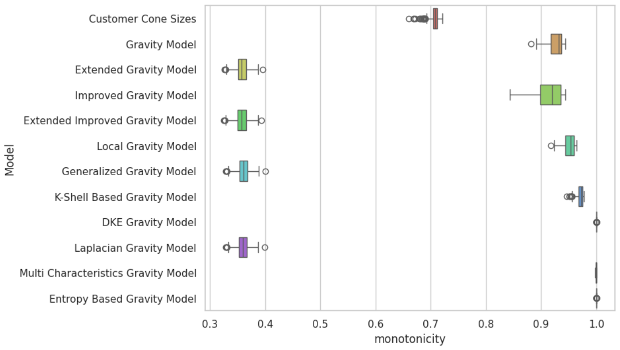

Fig. 3. Monotonicity scores across all snapshots **without leaf nodes**. Models like DKEGM, Entropy, and Multi-Characteristics achieve near-perfect ranking separation.

5.2 Rank Correlation

To understand how similar the rankings are, we calculated Kendall's τ and Spearman's ρ between every pair of models. For each metric, we computed the median correlation value across all CAIDA snapshots from 1998 to 2025. This ensures that short-term fluctuations do not dominate the results.

With leafs: In the Kendall and Spearman heatmaps (Figures 4 and 5) for rankings with leafs included, two clusters stand out. One contains models Gravity, Improved, Local, K-Shell, DKE, Multi, Entropy and Customer-Cone. The other groups Extended, Extended Improved, Laplacian, and Generalized Gravity.

While these clusters appear clear, they don't provide much insight. The main issue is that most leaf nodes receive the same rank (often 0), so all models share a large block of similar values. This drives the correlation scores up, but it does not reflect real structural similarity between the models. So although it looks structured, these correlations are skewed by the uniform treatment of leaves.

Without leafs: After removing leaf nodes, the heatmaps (Figures 6 and 7) become much more informative. The gravity-based models no longer form one big cluster instead, three groups emerge:

- The first includes Extended, Extended Improved, Laplacian, and Generalized Gravity. These share a specific design and correlate well.
- The second cluster is only the Multi-Characteristics Gravity model. Interestingly, it doesn't strongly correlate with any of the others even the gravity-based ones suggesting it captures a different structure.
- The third cluster includes Gravity, Improved, Local, K-Shell, DKE, and Entropy Gravity. These models rank ASes similarly.

The Customer-Cone metric is a clear outlier. It shows only a weak correlation with a few models, mostly Local Gravity, likely because both use some form of local reachability. However, even this connection is limited, showing that cone size reflects a very different way of ranking ASes.

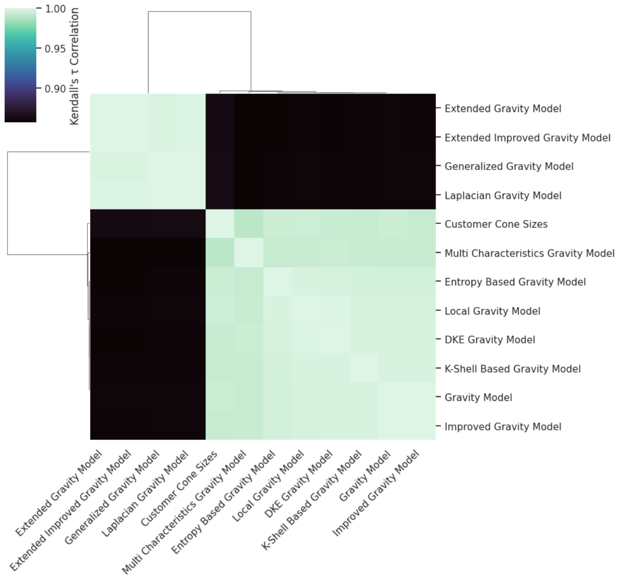

Fig. 4. Kendall rank correlation **with leaf nodes included**. Two clusters emerge, but the structure is inflated due to many identical leaf scores.

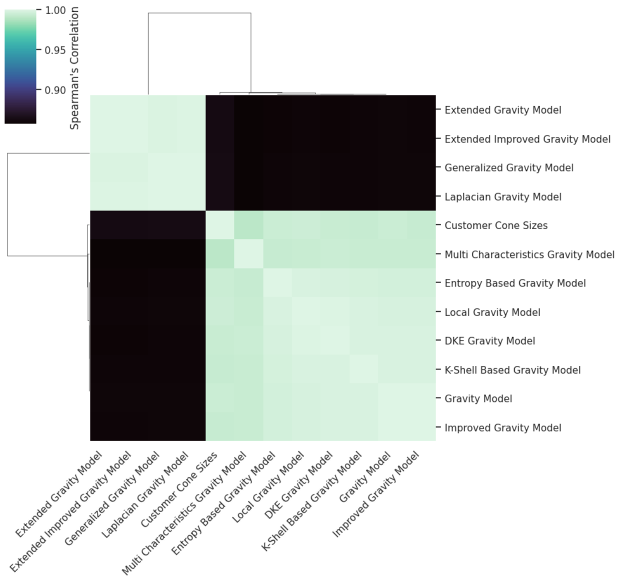

Fig. 5. Spearman rank correlation **with leaf nodes included**. Similar to Kendall: high correlations arise from uniformly zero leaf rankings.

5.3 Visual Ranking Comparison

Table 2 in Appendix A shows the top 10 ASes per model based on the CAIDA snapshot from May 1, 2025. Some ASes like *Level 3 Parent* and *Cogent* appear at the top across all models, which makes sense given their global connectivity. But there are small differences between methods.

The customer-cone ranking mostly favors ASes with large downstream trees. Gravity models, on the other hand, sometimes bump up ASes that are more central in structure, even if they do not have the biggest cones. This supports the idea that gravity-based models can offer more nuance, especially in hierarchical AS networks.

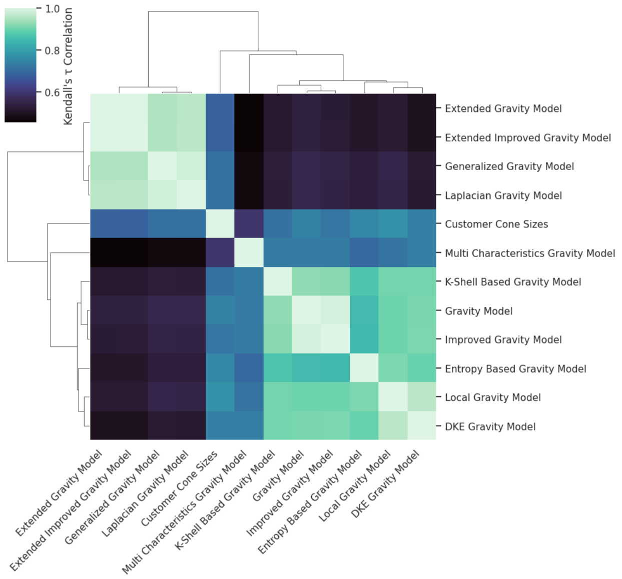

Fig. 6. Kendall rank correlation **without leaf nodes**. Three distinct clusters appear: (1) Extended models, (2) Multi-Characteristics (outlier), and (3) Core gravity models. Customer cone is weakly correlated only with Local Gravity.

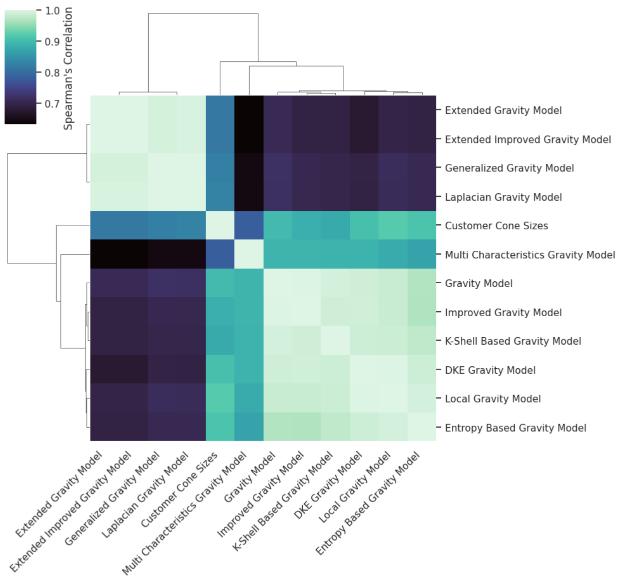

Fig. 7. Spearman rank correlation **without leaf nodes**. Cluster structure is preserved. Multi-Characteristics remains isolated; cone size still aligns only with Local Gravity.

6 DISCUSSION

6.1 Interpretability of Rankings

While some models (like DKE and Entropy-Based) perform very well in monotonicity and correlation, their internal mechanisms combine several metrics in a non-transparent way. For real-world applications, especially in network security or policymaking, understanding why a certain AS ranks high might be just as important as the score itself. Models like the basic Gravity or Local Gravity are easier to interpret and may be preferred in such settings.

6.2 Robustness Across Time

Given that our dataset spans from 1998 to 2025, we've implicitly tested whether models remain consistent across different AS graph structures and topologies. The strong performance of certain gravity models across all snapshots suggests they are more robust to structural changes over time. This is particularly valuable for long-term planning or detecting shifts in network influence.

6.3 Relevance to Real-World AS Properties

The results show that the models consistently highlight ASes such as Level 3, Cogent, and Arelion, which are well-known large Tier 1 providers. This supports that gravity-based models are capable of capturing meaningful structural properties in the AS topology. In contrast, customer-cone size may fail to identify ASes that are structurally important but have fewer direct customers, especially when their centrality arises from indirect or strategic connectivity.

Sprint ranks highly in some gravity-based models but only 2501st by customer-cone size. While it doesn't have a big customer cone, it plays an important role through indirect paths, such as: `Level3` $\longrightarrow$ `Yale University` $\longleftarrow$ `Sprint`. Such connections are not found by the customer cone but are visible in gravity-based rankings.

6.4 How Do Gravity Models Compare?

The monotonicity and rank correlation results show a clear separation between the stronger and weaker models. Gravity, Improved Gravity, Local Gravity, K-Shell, DKEGM, Multi-Characteristics, and Entropy-Based models all score highly in monotonicity especially after leaf nodes are removed and are consistently grouped in correlation heatmaps. However, a high rank correlation score does not always mean that a ranking is better. It just means that two models agree with each other. So, while correlation tells us how stable models are with each other, it doesn't tell us if they are correct. That's why we look at monotonicity alongside it, because it gives a better sense of how well a model separates ASes based on the topology.

6.5 Are Gravity-Based Rankings a Better Alternative?

Compared to traditional metrics like customer-cone size or degree, gravity-based models offer a more nuanced picture. They consider not just how many neighbors a node has, but also how important those neighbors are and how close they are in the network. Some also integrate additional properties like eigenvector centrality or entropy. This makes them more scalable and discriminative in large AS graphs, especially DAGs where plain degree or cone count might miss central ASes. Additionally, gravity scores work even if the graph changes slightly, making them more robust to topology shifts.

6.6 Relationship Between Gravity Centrality and Customer Cone Size

One interesting observation is that the Local Gravity model shows some correlation with Customer Cone Size probably because it also reflects local reach. This suggests it could be a good middle ground for operators who want something familiar but better. While customer cone focuses only on downstream count, gravity-based models add structural awareness. So, combining them might lead to an even more accurate ranking system that better reflects real-world AS influence, not just how many customers an AS has.

6.7 Limitations

Most gravity-based centrality models assume clean hierarchical propagation of influence, which doesn't always hold in real AS graphs due to peering links and multihoming. Also, metrics like entropy and eigenvector centrality rely on undirected assumptions or smooth connectivity, which may distort scores in a DAG setting. Another issue is how we handled leaves removing them improves metric quality, but also removes many smaller ASes, potentially biasing the analysis toward large Tier 1/2 providers.

7 CONCLUSION AND FUTURE WORK

7.1 Conclusion

This study set out to evaluate the suitability of gravity-based centrality models for ranking nodes in DAGs of Autonomous System (AS) networks. We posed three research questions: (1) How do gravity-based models perform in terms of ranking resolution and stability? (2) Can these models offer a more scalable and discriminative alternative to traditional metrics? (3) Can gravity centrality be used to enhance or complement customer-cone rankings?

We compared twelve gravity-based centrality variants against the widely used customer-cone size metric, using AS relationship snapshots from 1998 to 2025. Our evaluation focused on monotonicity and rank correlation, both with and without preprocessing steps such as leaf removal.

The results show that DKEGM, Entropy-Based Gravity, and Multi-Characteristics Gravity consistently produce clean, fine-grained rankings with near-perfect monotonicity, thereby answering RQ1 regarding resolution and stability. Specifically, without leaf nodes, DKEGM and Entropy-Based Gravity achieve monotonicity scores around 0.9999, while Multi-Characteristics Gravity closely follows with 0.9981, all substantially outperforming the customer-cone metric (0.7035). Even with leaves present, these models retain top performance, with scores around 0.0848, compared to 0.0825 for the cone.

In addressing RQ2, gravity-based models especially DKEGM, Entropy-Based, and Multi-Characteristics, demonstrate significantly higher discriminative power than the customer-cone size. DKEGM, for example, shows the strongest internal consistency with Kendall $\tau \approx 0.994$ and Spearman $\rho \approx 0.987$, while Multi-Characteristics and Entropy-Based models also exhibit high agreement across rankings, confirming their robustness and scalability.

Regarding RQ3, Local Gravity provides a middle ground: it aligns moderately with customer-cone rankings (Kendall $\tau = 0.769$ without leaves) while offering better resolution (monotonicity ≈ 0.9575), indicating its value as a complementary method for cone-based thinking.

Overall, gravity-based models offer a robust, scalable, and nuanced approach to AS centrality in DAG-structured networks. In practice, the choice of model depends on the intended application: for compatibility with cone-based reasoning and interpretability, Local Gravity is a natural fit due to its intuitive design and structural alignment. For applications requiring precise, stable, and fine-grained rankings, models such as DKEGM, Entropy-Based Gravity, and Multi-Characteristics Gravity are preferable, offering strong performance across all evaluation metrics while capturing subtle differences in network influence.

7.2 Future Work

There are several directions worth exploring:

- Track how rankings evolve over time to identify rising or declining ASes.
- Combine structural models with actual traffic data to validate rankings.
- Investigate how centralities changes after node removals.
- Some gravity models have tunable parts (like distance functions). Fine-tuning these may yield better results.
- Whether results hold across other different DAG-structured networks, not just AS graphs.

ACKNOWLEDGEMENTS

I'd like to thank my supervisors, Dr. Zangiabady and Dr. Garcia-Robledo, for their help and support throughout this project. Their feedback and advice really helped me stay on track and improve my work. I've learned a lot from their guidance and I'm glad I got to work with them.

REFERENCES

[1] Phillip Bonacich and Paulette Lloyd. 2001. Eigenvector-like measures of centrality for asymmetric relations. *Social Networks* 23, 3 (2001), 191–201. doi:10.1016/S0378-8733(01)00038-7

[2] CAIDA (The Center for Applied Internet Data Analysis). 2024. CAIDA AS Relationships Dataset. https://www.caida.org/catalog/datasets/as-relationships/. Accessed: 2025-05-21.

[3] Cloudflare. [n. d.]. What is an autonomous system? | What are ASNs? | Cloudflare. *Cloudflare* ([n. d.]). https://www.cloudflare.com/learning/network-layer/what-is-an-autonomous-system/

[4] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking* 9, 6 (2001), 733–745. doi:10.1109/90.974527

[5] J. Hawkinson and T. Bates. 1996. RFC1930: Guidelines for creation, selection, and registration of an Autonomous System (AS).

[6] Amy N. Langville and Carl D. Meyer. 2006. *Google's PageRank and Beyond: The Science of Search Engine Rankings.* Princeton University Press. http://www.jstor.org/stable/j.ctt7t8z9

[7] Hanwen Li, Qiuyan Shang, and Yong Deng. 2021. A generalized gravity model for influential spreaders identification in complex networks. *Chaos, Solitons & Fractals* 143 (2021), 110456. doi:10.1016/j.chaos.2020.110456

[8] Zhe Li and Xinyu Huang. 2022. Identifying influential spreaders by gravity model considering multi-characteristics of nodes. *Scientific Reports* 12, 1 (Jun 2022). doi:10.1038/s41598-022-14005-3

[9] Zhe Li, Tao Ren, Xiaoqi Ma, Simiao Liu, Yixin Zhang, and Tao Zhou. 2019. Identifying influential spreaders by gravity model. *Scientific Reports* 9, 1 (Jun 2019). doi:10.1038/s41598-019-44930-9

[10] Ling ling Ma, Chuang Ma, Hai-Feng Zhang, and Bing-Hong Wang. 2016. Identifying influential spreaders in complex networks based on gravity formula. *Physica A: Statistical Mechanics and its Applications* 451 (2016), 205–212. doi:10.1016/j.physa.2015.12.162

[11] Yong Liu, Zijun Cheng, Xiaoqin Li, and Zongshui Wang. 2023. An Entropy-Based Gravity Model for Influential Spreaders Identification in Complex Networks. *Complexity* 2023, 1 (2023), 6985650. doi:10.1155/2023/6985650 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1155/2023/6985650

[12] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy. 2013. AS relationships, customer cones, and validation. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (Barcelona, Spain) *(IMC '13)*. Association for Computing Machinery, New York, NY, USA, 243–256. doi:10.1145/2504730.2504735

[13] Tore Opsahl, Filip Agneessens, and John Skvoretz. 2010. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks* 32, 3 (2010), 245–251. doi:10.1016/j.socnet.2010.03.006

[14] Stephen B. Seidman. 1983. Network structure and minimum degree. *Social Networks* 5, 3 (1983), 269–287. doi:10.1016/0378-8733(83)90028-X

[15] Hardeep Singh and Hardeep Singh. 2024. An Enhanced Gravity Model for Determining Crucial Nodes in Social Networks Based on Degree K-Shell Eigenvector Index. *IEEE Access* 12 (2024), 23163–23180. doi:10.1109/ACCESS.2024.3363635

[16] Mehmet Engin Tozal. 2017. Autonomous system ranking by topological characteristics: A comparative study. In *2017 Annual IEEE International Systems Conference (SysCon)*. 1–8. doi:10.1109/SYSCON.2017.7934814

[17] Juan Wang, Chao Li, and Chengyi Xia. 2018. Improved centrality indicators to characterize the nodal spreading capability in complex networks. *Appl. Math. Comput.* 334 (2018), 388–400. doi:10.1016/j.amc.2018.04.028

[18] Xuan Yang and Fuyuan Xiao. 2021. An improved gravity model to identify influential nodes in complex networks based on k-shell method. *Knowledge-Based Systems* 227 (2021), 107198. doi:10.1016/j.knosys.2021.107198

[19] Qinyu Zhang, Bin Shuai, and Min Lü. 2022. A novel method to identify influential nodes in complex networks based on gravity centrality. *Information Sciences* 618 (2022), 98–117. doi:10.1016/j.ins.2022.10.070

AI STATEMENT

During the preparation of this work the author used ChatGPT in order to enhance the writing of this paper. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the work.

A FULL TOP-10 AS RANKINGS

Rank	Customer Cone	Gravity	Ext. Gravity	Imp. Gravity	Ext. Imp. Gravity	Local Gravity	Gen. Gravity	K-Shell Gravity	DKE Gravity	Laplacian Gravity	Multi-Char. Gravity	Entropy Gravity
1	Level 3 Parent, LLC	Cogent Comm.	Level 3 Parent, LLC	Cogent Comm.	Level 3 Parent, LLC	Level 3 Parent, LLC	Level 3 Parent, LLC	Level 3 Parent, LLC	Level 3 Parent, LLC	Level 3 Parent, LLC	Cogent Comm.	Cogent Comm.
2	Arelion Sweden AB	Level 3 Parent, LLC	Arelion Sweden AB	Level 3 Parent, LLC	Arelion Sweden AB	Cogent Comm.	Cogent Comm.	Cogent Comm.	Cogent Comm.	Cogent Comm.	Level 3 Parent, LLC	Level 3 Parent, LLC
3	Cogent Comm.	Arelion Sweden AB	Cogent Comm.	Arelion Sweden AB	Cogent Comm.	Arelion Sweden AB	Arelion Sweden AB	Arelion Sweden AB	Arelion Sweden AB	Arelion Sweden AB	Sprint	Sprint
4	GTT Comm. Inc.	GTT Comm. Inc.	NTT America, Inc.	GTT Comm. Inc.	TATA Comm.	GTT Comm. Inc.	GTT Comm. Inc.	GTT Comm. Inc.	Hurricane Electric LLC	GTT Comm. Inc.	Arelion Sweden AB	Zayo Bandwidth
5	Verizon Business	NTT America, Inc.	TATA Comm.	NTT America, Inc.	NTT America, Inc.	Hurricane Electric LLC	Hurricane Electric LLC	NTT America, Inc.	GTT Comm. Inc.	Hurricane Electric LLC	GTT Comm. Inc.	AT&T Services, Inc.
6	PCCW Global, Inc.	Hurricane Electric LLC	Telecom Italia	Hurricane Electric LLC	Telecom Italia	NTT America, Inc.	NTT America, Inc.	Hurricane Electric LLC	NTT America, Inc.	Zayo Bandwidth	Hurricane Electric LLC	Verizon Business
7	Hurricane Electric LLC	Telecom Italia	GTT Comm. Inc.	Telecom Italia	GTT Comm. Inc.	Zayo Bandwidth	Zayo Bandwidth	Zayo Bandwidth	Zayo Bandwidth	NTT America, Inc.	Zayo Bandwidth	Arelion Sweden AB
8	NTT America, Inc.	TATA Comm.	PCCW Global	TATA Comm.	PCCW Global	Verizon Business	Verizon Business	Verizon Business	Sprint	Verizon Business	NTT America, Inc.	GTT Comm. Inc.
9	Telecom Italia	PCCW Global	Durand do Brasil	PCCW Global	GTT Comm. Inc.	Sprint	Telecom Italia	PCCW Global	Verizon Business	PCCW Global	Verizon Business	Hurricane Electric LLC
10	TATA Comm.	Zayo Bandwidth	Hurricane Electric LLC	Durand do Brasil	Durand do Brasil	Telecom Italia	RETN Limited	Telecom Italia	RETN Limited	RETN Limited	AT&T Services, Inc.	PJSC Rostelecom

Table 2. Top-10 AS rankings based on the May 1, 2025 CAIDA snapshot. Gravity-based models show strong agreement on core ASes, but differences remain, e.g., Sprint ranks high under Entropy, but very low by cone size.