





Metric Space Magnitude for Evaluating the Diversity of Latent Representations

Katharina Limbeck, Rayna Andreeva, Rik Sarkar, Bastian Rieck






Why do we want to evaluate diversity?

data / task

automated
assessment of
e.g. generative
model outputs



latent spaces




diversity evaluation

intrinsic diversity
reference-based
comparison

How to measure diversity?

Diversity measures the **effective size / clusterability / entropy** of a space.


X_1 is more diverse than X_2 . Baseline measures fail to detect this!



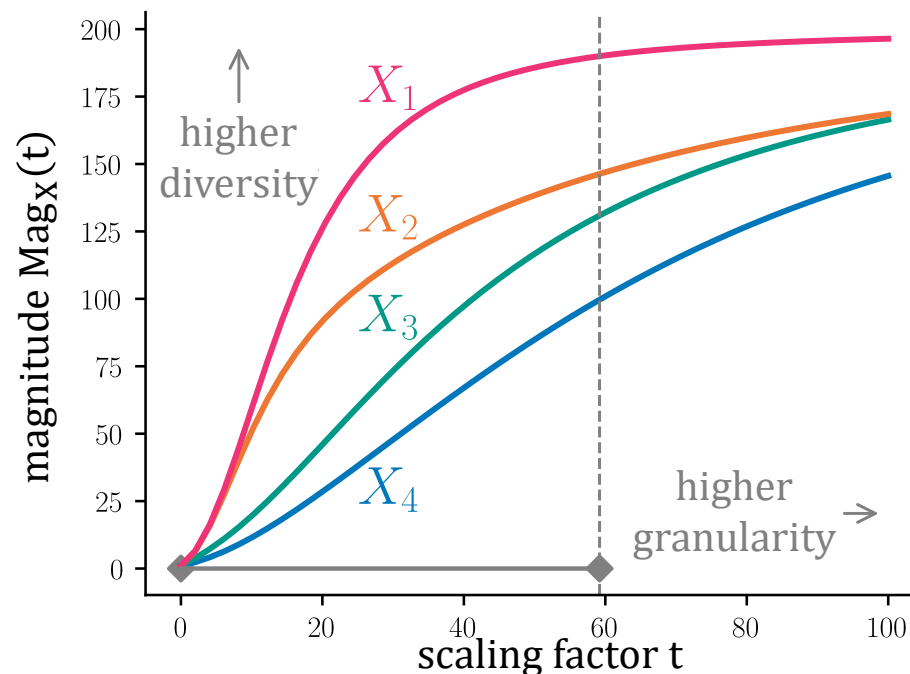


Magnitude magnifies diversity

Magnitude functions measure the *effective number of distinct points* at a scale of distance t .



latent representations




Contributions




MagArea, the area under the magnitude function, to measure *intrinsic diversity* of latent spaces.

MagDiff, the area between two magnitude functions, to measure the *difference in diversity*.


Automated computation and scale-selection.



Results



Magnitude detects **mode collapse and mode dropping** for graph and image generative models.



Magnitude **distinguishes text embedding models** and characterises the intrinsic diversity of generated text.



Magnitude for measuring diversity

Magnitude gives **theoretically justified, automated and multi-scale diversity measures** for evaluating latent representations.

