CT-DegradBench: A Physics-Informed Benchmark for CT Degradation Detection and Severity Estimation

Supplementary Material

6. Additional Quantitative Evaluation

6.1. Per-Degradation Classification Performance

Table 5 reports classification accuracy and F1-score for each degradation across single distortions (S1–S5) and mixtures (M1–M5). Single degradations achieve consistently high performance, with most classes exceeding 0.98 accuracy. Blur (S2) and metal artifacts (S5) obtain the highest F1-scores (0.960 and 0.961), indicating that the proposed method effectively captures their distinctive spatial and spectral patterns.

In contrast, noise and aliasing exhibit lower F1-scores (0.587 and 0.252, respectively). This behavior likely reflects the difficulty of modeling aliasing artifacts, which arise from angular undersampling and are primarily expressed as high-frequency sampling patterns that are less effectively encoded by representations learned from semantic visual features. As expected, mixture degradations are generally more challenging than single distortions, since overlapping artifacts may partially obscure individual degradation cues. In particular, M4 (aliasing + noise) yields the lowest F1-score, suggesting increased ambiguity when frequency-domain artifacts interact with noise.

Single Distortions			Mixture Distortions		
Degradation	Accuracy ↑	F1 ↑	Degradation	Accuracy ↑	F1 ↑
S1_noise	0.914	0.587	M1_b+n	0.935	0.598
S2_blur	0.992	0.960	M2_s+n	0.893	0.618
S3_streak_nm	0.982	0.911	M3_m+n	0.943	0.758
S4_aliasing	0.914	0.252	M4_a+n	0.876	0.438
S5_metal	0.992	0.961	M5_m+b+n	0.955	0.757

Mixture shorthand: b = blur, n = noise, s = streaks, a = aliasing, m = metal.
↑ higher is better.

Table 5. **Per-degradation classification performance of the proposed method.** One-vs-rest accuracy and F1-score for single distortions (S1–S5) and mixtures (M1–M5).

6.2. ROC Analysis for Degradation Classification

Figure 6 presents ROC curves for all degradations. The micro-average AUC reaches 0.954, indicating strong overall degradations discrimination. Single degradations achieve near-perfect AUC values in several cases, confirming that the learned representation clearly separates artifact types. Although mixture degradations show slightly lower AUC values, their performance remains high, demonstrating that the model retains discriminative features even under combined artifacts.

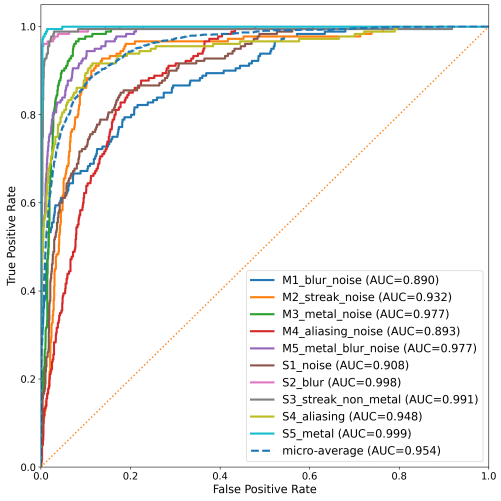

Figure 6. **ROC curves for degradation classification.** One-vs-rest ROC curves are shown for each degradation category. The dashed line denotes the micro-average across all classes (AUC = 0.954).

6.3. Severity Level Prediction Performance

Table 6 reports severity prediction performance using MAE, RMSE, and Quadratic Weighted Kappa (QWK). For single degradations, blur and streak artifacts achieve the lowest errors. Metal artifacts are more challenging to estimate, likely due to their dependence on implant geometry and position. Mixtures show moderate prediction errors but maintain reasonable agreement with ground truth, indicating that the model captures severity-related cues even in the presence of interacting degradations.

Single Distortions				Mixture Distortions			
Degradation	MAE ↓	RMSE ↓	QWK ↑	Degradation	MAE ↓	RMSE ↓	QWK ↑
S1_noise	0.578	0.713	0.754	M1_b+n	0.821	1.030	0.656
S2_blur	0.387	0.567	0.869	M2_s+n	0.785	0.994	0.601
S3_streak_nm	0.495	0.626	0.803	M3_m+n	0.623	0.794	0.699
S4_aliasing	0.659	0.808	0.772	M4_a+n	0.820	1.009	0.566
S5_metal	0.989	1.242	0.396	M5_m+b+n	0.599	0.771	0.716

Mixture shorthand: b = blur, n = noise, s = streaks, a = aliasing, m = metal.
↑ higher is better; ↓ lower is better.

Table 6. **Per-degradation severity prediction performance.** MAE, RMSE, and Quadratic Weighted Kappa (QWK) for severity estimation across single distortions (S1–S5) and mixtures (M1–M5).

7. Embedding-Space Representation Analysis

7.1. t-SNE Visualization by Degradation Type

Figure 5 visualizes the learned embedding space projected using t-SNE and colored by degradation category. Distinct clusters emerge for most degradation types, indicating that the proposed representation encodes artifact-specific characteristics. Degradations that introduce strong structural patterns, such as blur and metal artifacts, form compact and well-separated clusters, suggesting that their spatial and spectral signatures are consistently captured by the model.

In contrast, degradations dominated by stochastic or high-frequency components, such as noise and aliasing, exhibit slightly more diffuse distributions. This behavior is expected, as these artifacts may share overlapping frequency characteristics with other degradations. For example, samples affected by aliasing occasionally appear closer to streak or noise clusters, reflecting the partial similarity between undersampling patterns and other frequency-domain distortions.

7.2. t-SNE Visualization by Severity Level

Figure 7 shows the same embedding space colored by degradation severity. A gradual transition between severity levels can be observed within several clusters, indicating that the representation captures not only degradation type but also variations in degradation intensity. In particular, blur and streak artifacts display a clear progression from lower to higher severity levels along consistent directions in the embedding space. entirely separate regions. However,

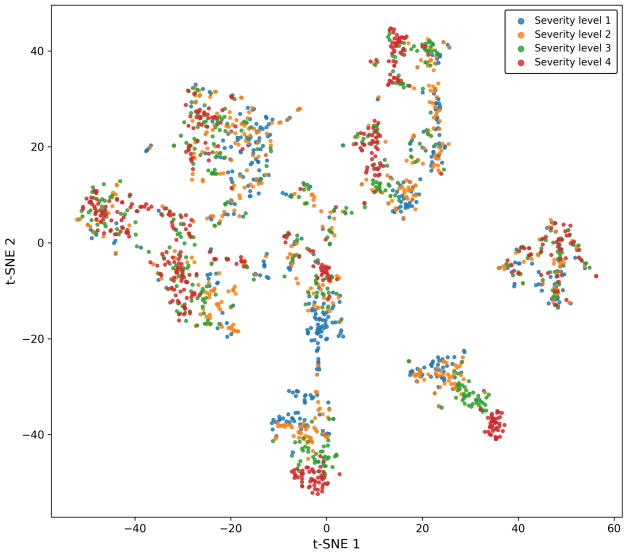

Figure 7. **t-SNE visualization colored by degradation severity.** The embedding reveals a gradual structure where samples are organized according to increasing degradation intensity.

severity separation is less pronounced for noise and metal artifacts. This emphasizes the results presented previously in tables 5 and 6 for these degradations, where severity do not severely impact the visual characteristics.

7.3. t-SNE Visualization for Single and Mixed Degradations

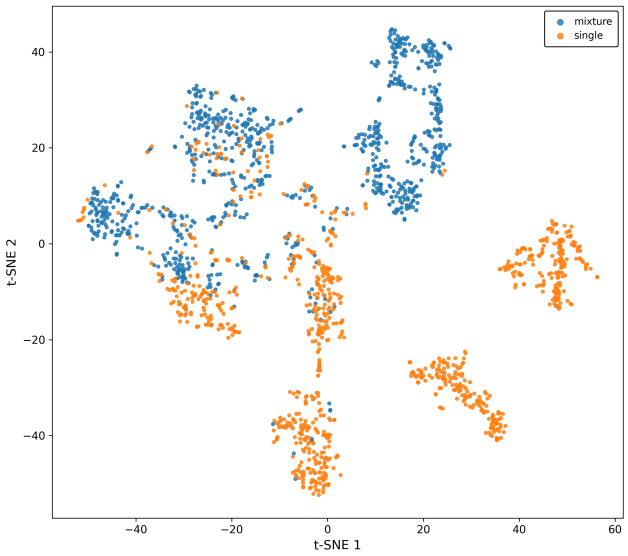

Figure 8. **t-SNE visualization distinguishing single and mixed degradations.** The representation separates samples containing a single degradation from those with multiple degradations, highlighting its ability to capture compositional corruption patterns.

Figure 8 highlights the distribution of single degradations versus mixtures. While many mixture samples remain close to the clusters associated with their dominant artifact type, they often appear near the boundaries between multiple clusters. This behavior suggests that mixed degradations combine characteristics from several artifact types, resulting in embeddings that lie between the corresponding single-degradation regions. For example, mixtures involving metal artifacts and noise tend to appear near the metal artifact cluster while exhibiting slight displacement toward regions associated with noise. These patterns indicate that the learned embedding space captures compositional degradation characteristics rather than assigning mixtures to

8. Extended Ablation Studies

8.1. Branch Contribution Analysis

Table 4 evaluates the contribution of each branch in the proposed semantic-spectral framework. Ablation results reveal that the two branches play complementary roles in capturing degradation characteristics.

Removing the semantic quality branch (B_{sem}) causes a drastic drop in classification performance (Acc = 0.151, F1

= 0.061), indicating that the semantic representation learned from the vision–language model provides a crucial prior for distinguishing degradation types. Without this branch, the model relies primarily on frequency information, which alone shows insufficient performance to discriminate between artifacts.

In contrast, removing the frequency branch (B_{FFT}) significantly reduces performance but to a lesser extent (Acc = 0.698, F1 = 0.684). This suggests that while semantic representations capture high-level structural cues, the spectral branch provides complementary information about high-frequency artifact patterns. Degradations such as aliasing, streak artifacts, and noise are particularly characterized by distinct frequency signatures, which explains the notable performance gain when the spectral branch is included.

8.2. Loss Component Analysis

The ablation results also highlight the importance of the different training objectives. Removing the regression loss (L_{reg}) leads to a substantial degradation in severity estimation performance (RMSE increases from 0.794 to 1.783), confirming that explicit regression supervision is essential for learning a continuous severity representation.

Similarly, removing the ranking loss (L_{rank}) noticeably affects the severity prediction metrics. The increase in RMSE and decrease in QWK suggest that the ranking constraint plays an important role in preserving the ordinal structure of degradation severity levels. In other words, it encourages the model to maintain consistent ordering between different severity levels, which is particularly important when the visual differences between adjacent levels are subtle.

Finally, removing the contrastive loss (L_{con}) results in a decrease in both classification and severity estimation performance. This observation indicates that contrastive supervision helps structure the embedding space by pulling together samples with similar degradation characteristics while pushing apart dissimilar ones.

Overall, the ablation study demonstrates that the semantic and spectral branches provide complementary information, while the combination of regression, ranking, and contrastive objectives encourages a representation space that is both discriminative and consistent with degradation severity.

9. Additional Benchmark Details

This section provides supplementary details on the prompt design used to construct the semantic quality axis, the physical motivation behind the degradation mixtures in CT-DegradBench, and the relation of CT-DegradBench to existing CT restoration datasets.

9.1. Prompt design for the semantic quality axis

Table 7 lists the high- and low-quality prompts used to construct the semantic quality axis in Section 3. The high-quality prompts describe diagnostically reliable CT images with clear anatomy, sharp boundaries, and no visible artifacts, whereas the low-quality prompts describe degraded CT appearances corresponding to the main artifact families in CT-DegradBench: noise, blur, streaks, aliasing, and metal artifacts. These prompts define the semantic direction used by the semantic branch without task-specific fine-tuning.

9.2. Clinical motivation for degradation mixtures

Table 8 summarizes the motivating scenarios and physical interpretation of the degradation mixtures used in CT-DegradBench. These mixtures are designed to reflect plausible CT acquisition conditions in which multiple degradation sources co-occur, particularly under low-dose imaging, sparse-view acquisition, or the presence of metallic implants.

9.3. Comparison with existing CT datasets

Table 9 compares CT-DegradBench with commonly used CT restoration datasets. Most existing datasets focus on a single degradation family, such as low-dose noise, sparse-view artifacts, or metal artifact reduction, and typically lack explicit severity control or mixed-degradation settings. In contrast, CT-DegradBench provides multiple degradation types, calibrated severity levels, and controlled mixtures within a unified benchmark.

High-Quality Prompts (H)	Low-Quality Prompts (L)
1. Axial abdominal CT slice with excellent diagnostic quality, sharp boundaries, clear organ detail, and no visible artifacts.	1. Abdominal CT slice with severe noise and grainy appearance that reduces visibility of anatomical structures.
2. Diagnostic abdominal CT with clear anatomical structures, low noise, high contrast, and no streak artifacts.	2. Abdominal CT slice with strong blur and significant loss of sharpness.
3. High-quality CT image with sharp edges, clean appearance, and good visibility of abdominal organs.	3. Abdominal CT slice with strong streak artifacts and reduced diagnostic quality.
	4. Abdominal CT slice with sparse-view aliasing artifacts and distorted anatomical structures.
	5. Abdominal CT slice with strong metal artifacts causing bright streaks and severe image corruption.

Table 7. **Prompt sets used to construct the semantic quality axis.** High-quality prompts describe artifact-free diagnostic CT images, while low-quality prompts represent common degradations including noise, blur, streak artifacts, aliasing, and metal artifacts.

Mixture	Description
Loss of sharpness + Noise	*Scenario:* Low-dose CT with limited detector resolution. *Mechanism:* Detector response (MTF) reduces spatial resolution (loss of sharpness) → low photon counts introduce noise.
Streaks + Noise	*Scenario:* Photon starvation in dense regions (e.g., shoulders or pelvis). *Mechanism:* Highly attenuating paths produce streak artifacts → low-dose acquisition increases noise.
Metal + Noise	*Scenario:* CT scans acquired at low tube current with metallic implants (e.g., hip prosthesis). *Mechanism:* Photon starvation from metal produces streak artifacts → low photon counts introduce noise.
Aliasing + Noise	*Scenario:* Accelerated or sparse-view low-dose CT acquisition. *Mechanism:* Angular undersampling produces aliasing → reduced photon counts introduce noise.
Metal + Loss of sharpness + Noise	*Scenario:* CT imaging with metallic implants under low-dose conditions. *Mechanism:* Metal beam hardening → detector response reduces spatial resolution → low photon counts introduce noise.

Table 8. Scenario motivation and physical interpretation of degradation mixtures in CT-DegradBench.

Dataset	Task	Degradation Types	Severity per Degradation	Severity Control	Mixtures
AAPM Mayo LDCT [13]	LDCT enhancement	Low-dose photon noise	Quarter dose level	✗	Implicit
LDCT-and-Projection-data [15]	LDCT enhancement	Low-dose degradations	Quarter dose level	✗	Implicit
LoDoPaB-CT [12]	LDCT enhancement	Low-dose degradations	Quarter dose level	✗	Implicit
Piglet LDCT Dataset [27]	LDCT enhancement	Low-dose degradations	Five dose levels	✓	Implicit
AAPM Sparse-View CT Challenge [21]	Sparse-view reconstruction	Angular undersampling	Single level (128 views)	✗	✗
AAPM CT Metal Artifact Reduction Challenge [8]	Metal Artifact Reduction (MAR)	Simulated metal streak artifact	Not graded	✗	✗
UCLH Stroke EIT Dataset [6]	Metal Artifact Reduction (MAR)	Metal-induced streak artifacts	Not explicitly graded	✗	✗
CT-DegradBench (Ours)	Degradation benchmarking	Noise, blur, streaks, aliasing, metal artifacts	**4 levels**	✓	✓

Table 9. Comparison with common CT datasets used in restoration and enhancement pipelines.

Table 10. Benchmarking results (mean ± std) for full-reference quality metrics across degradations and severity levels.

Single distortions						Mixtures					
Setting	**PSNR ↑**	**SSIM ↑**	**VIF ↑**	**LPIPS ↓**	**DISTS ↓**	**Setting**	**PSNR ↑**	**SSIM ↑**	**VIF ↑**	**LPIPS ↓**	**DISTS ↓**
S1_noise L0	29.67±1.95	0.8627±0.0482	0.3657±0.0877	0.3650±0.0730	0.1917±0.0505	M1_b+n L0	28.20±2.31	0.7931±0.0909	0.2586±0.0615	0.4400±0.0698	0.2401±0.0482
S1_noise L1	27.14±2.44	0.7456±0.0978	0.2695±0.0787	0.4564±0.0730	0.2569±0.0484	M1_b+n L1	27.54±2.56	0.7600±0.1089	0.2372±0.0647	0.4599±0.0741	0.2549±0.0503
S1_noise L2	25.96±2.68	0.6920±0.1142	0.2403±0.0739	0.4839±0.0710	0.2766±0.0453	M1_b+n L2	25.56±2.85	0.6650±0.1273	0.1813±0.0557	0.5149±0.0634	0.2936±0.0400
S1_noise L3	23.08±3.10	0.5638±0.1373	0.1827±0.0631	0.5383±0.0657	0.3135±0.0384	M1_b+n L3	24.40±2.77	0.6057±0.1224	0.1417±0.0380	0.5452±0.0520	0.3157±0.0298
S2_blur L0	30.96±1.80	0.9260±0.0108	0.4629±0.0431	0.3539±0.0280	0.1363±0.0164	M2_s+n L0	28.44±2.44	0.8027±0.0934	0.2973±0.0853	0.4254±0.0777	0.2319±0.0541
S2_blur L1	30.83±1.74	0.9225±0.0107	0.4264±0.0402	0.3787±0.0285	0.1525±0.0156	M2_s+n L1	27.77±2.70	0.7720±0.1080	0.2769±0.0826	0.4453±0.0766	0.2455±0.0527
S2_blur L2	30.49±1.60	0.9115±0.0104	0.3534±0.0326	0.4128±0.0264	0.1881±0.0127	M2_s+n L2	24.92±2.99	0.6505±0.1285	0.2109±0.0670	0.5094±0.0670	0.2881±0.0434
S2_blur L3	29.81±1.37	0.8846±0.0114	0.2541±0.0231	0.4600±0.0232	0.2370±0.0110	M2_s+n L3	22.92±2.81	0.5642±0.1235	0.1744±0.0562	0.5442±0.0578	0.3119±0.0361
S3_streak L0	31.58±1.69	0.9385±0.0082	0.5799±0.0444	0.2136±0.0145	0.0987±0.0130	M3_m+n L0	23.89±0.90	0.7090±0.0622	0.2046±0.0406	0.4977±0.0406	0.2844±0.0284
S3_streak L1	31.28±1.59	0.9256±0.0094	0.5575±0.0439	0.2485±0.0165	0.1233±0.0165	M3_m+n L1	23.05±1.45	0.6425±0.1087	0.1818±0.0503	0.5221±0.0537	0.3002±0.0351
S3_streak L2	30.29±1.31	0.8936±0.0166	0.5187±0.0439	0.2953±0.0180	0.1593±0.0197	M3_m+n L2	20.89±1.66	0.4884±0.1067	0.1296±0.0357	0.5763±0.0426	0.3369±0.0233
S3_streak L3	27.83±0.87	0.8298±0.0305	0.4649±0.0440	0.3442±0.0199	0.1968±0.0207	M3_m+n L3	19.19±1.13	0.3833±0.0694	0.0996±0.0226	0.6065±0.0322	0.3567±0.0158
S4_alias L0	31.05±1.84	0.9234±0.0110	0.5232±0.0444	0.2343±0.0175	0.0923±0.0149	M4_a+n L0	28.06±2.28	0.7750±0.0819	0.2729±0.0641	0.4513±0.0589	0.2552±0.0367
S4_alias L1	30.33±1.59	0.8647±0.0172	0.3936±0.0314	0.3641±0.0187	0.2069±0.0173	M4_a+n L1	27.48±2.27	0.7439±0.0901	0.2543±0.0638	0.4680±0.0589	0.2663±0.0357
S4_alias L2	29.40±1.33	0.7896±0.0268	0.3076±0.0244	0.4322±0.0157	0.2538±0.0127	M4_a+n L2	25.25±2.60	0.6134±0.1036	0.1885±0.0523	0.5311±0.0515	0.3018±0.0271
S4_alias L3	28.57±1.16	0.7267±0.0331	0.2546±0.0202	0.4697±0.0138	0.2793±0.0101	M4_a+n L3	23.87±2.79	0.5478±0.1020	0.1606±0.0409	0.5560±0.0485	0.3149±0.0258
S5_metal L0	25.28±0.86	0.8527±0.0197	0.3780±0.0397	0.3940±0.0291	0.2004±0.0179	M5_m+b+n L0	24.11±0.94	0.7314±0.0676	0.1960±0.0385	0.4890±0.0425	0.2777±0.0293
S5_metal L1	24.61±0.80	0.8414±0.0204	0.3686±0.0385	0.3946±0.0279	0.2058±0.0177	M5_m+b+n L1	23.11±1.34	0.6469±0.1022	0.1638±0.0420	0.5262±0.0491	0.3023±0.0317
S5_metal L2	24.02±0.74	0.8311±0.0213	0.3597±0.0373	0.3960±0.0273	0.2109±0.0177	M5_m+b+n L2	21.13±1.52	0.4979±0.1023	0.1128±0.0306	0.5802±0.0387	0.3391±0.0217
S5_metal L3	23.58±0.69	0.8221±0.0222	0.3519±0.0363	0.3976±0.0265	0.2144±0.0174	M5_m+b+n L3	19.52±1.44	0.3942±0.0918	0.0804±0.0195	0.6112±0.0340	0.3595±0.0187

Setting keywords: S1-S5 = single degradations; M1-M5 = mixtures; L0-L3 = severity levels (increasing with L).
Mixture shorthand: b = blur, n = noise, s = streaks, a = aliasing, m = metal.
↑ higher is better; ↓ lower is better.

Single distortions (ρ / r)					Mixtures (ρ / r)				
Setting	**OpenCLIP [3]**	**MedCLIP [25]**	**BioMedCLIP [31]**	**Merlin [1]**	**Setting**	**OpenCLIP [3]**	**MedCLIP [25]**	**BioMedCLIP [31]**	**Merlin [1]**
S1_noise	0.6414/0.6422	0.3253/0.2931	0.6401/0.6046	0.6219/0.5272	M1_b+n	0.5453/0.5557	0.2470/0.2269	0.5963/0.5807	0.4139/0.3512
S2_blur	0.5395/0.5261	0.5527/0.5128	0.8854/0.7817	0.1712/0.1679	M2_s+n	0.6134/0.6120	0.2706/0.2622	0.6823/0.6613	0.5126/0.4367
S3_streak	0.7734/0.6647	0.5770/0.5311	0.8962/0.8428	0.2818/0.2862	M3_m+n	0.7462/0.7448	0.1985/0.1941	0.6272/0.6088	0.6717/0.5891
S4_aliasing	0.8876/0.8419	0.7371/0.6268	0.9115/0.8251	0.5651/0.5180	M4_a+n	0.7148/0.6939	0.2230/0.2113	0.6473/0.6036	0.4597/0.4152
S5_metal	0.2864/0.2845	0.0778/0.0688	0.0989/0.0957	0.0645/0.0640	M5_m+b+n	0.7411/0.7438	0.1906/0.1985	0.6569/0.6402	0.6679/0.5830
Mean	0.6257/0.5919	0.4540/0.4065	**0.6864/0.6300**	0.3409/0.3127	**Mean**	**0.6722/0.6700**	0.2259/0.2186	0.6420/0.6189	0.5452/0.4750

Entries are reported as ρ/r (Spearman/Pearson).
Setting keywords: S1-S5 = single degradations; M1-M5 = mixtures.

Table 11. Correlation between VLM embedding drift and degradation severity for single degradations and degradation mixtures. Entries are reported as Spearman/Pearson correlations. (Full Dataset)

Table 12. Benchmarking results (mean ± std) for VLM embedding drift across degradations and severity levels).

Single distortions					Mixtures				
Setting	**OpenCLIP [3]**	**MedCLIP [25]**	**BioMedCLIP [31]**	**Merlin [1]**	**Setting**	**OpenCLIP [3]**	**MedCLIP [25]**	**BioMedCLIP [31]**	**Merlin [1]**
S1_noise L0	0.0636±0.0255	0.0537±0.0471	0.0150±0.0121	0.0021±0.0024	M1_b+n L0	0.0885±0.0350	0.0732±0.0479	0.0475±0.0332	0.0049±0.0066
S1_noise L1	0.0968±0.0374	0.0767±0.0523	0.0503±0.0390	0.0069±0.0070	M1_b+n L1	0.1003±0.0415	0.0780±0.0497	0.0614±0.0474	0.0058±0.0065
S1_noise L2	0.1147±0.0420	0.0844±0.0543	0.0697±0.0510	0.0102±0.0096	M1_b+n L2	0.1382±0.0556	0.0933±0.0545	0.1087±0.0662	0.0111±0.0108
S1_noise L3	0.1722±0.0660	0.0999±0.0592	0.1258±0.0793	0.0205±0.0171	M1_b+n L3	0.1752±0.0613	0.1046±0.0580	0.1548±0.0736	0.0141±0.0137
S2_blur L0	0.0492±0.0169	0.0329±0.0199	0.0100±0.0056	0.0007±0.0005	M2_s+n L0	0.0837±0.0336	0.0700±0.0490	0.0491±0.0319	0.0041±0.0048
S2_blur L1	0.0519±0.0180	0.0395±0.0236	0.0144±0.0072	0.0007±0.0005	M2_s+n L1	0.0925±0.0373	0.0769±0.0514	0.0666±0.0488	0.0057±0.0076
S2_blur L2	0.0601±0.0211	0.0567±0.0324	0.0307±0.0131	0.0007±0.0005	M2_s+n L2	0.1405±0.0551	0.1002±0.0654	0.1297±0.0681	0.0125±0.0145
S2_blur L3	0.0872±0.0273	0.0796±0.0375	0.0729±0.0269	0.0009±0.0006	M2_s+n L3	0.1814±0.0620	0.1089±0.0630	0.1736±0.0624	0.0181±0.0147
S3_streak L0	0.0366±0.0140	0.0183±0.0127	0.0084±0.0066	0.0004±0.0003	M3_m+n L0	0.1884±0.0572	0.1046±0.0569	0.2022±0.0494	0.0052±0.0050
S3_streak L1	0.0451±0.0163	0.0296±0.0179	0.0331±0.0244	0.0004±0.0003	M3_m+n L1	0.2331±0.0908	0.1110±0.0593	0.2237±0.0505	0.0088±0.0083
S3_streak L2	0.0629±0.0253	0.0453±0.0272	0.0927±0.0415	0.0005±0.0004	M3_m+n L2	0.3337±0.0766	0.1271±0.0677	0.2657±0.0536	0.0185±0.0150
S3_streak L3	0.1029±0.0403	0.0580±0.0332	0.1568±0.0492	0.0007±0.0005	M3_m+n L3	0.3924±0.0508	0.1366±0.0692	0.3010±0.0430	0.0279±0.0155
S4_alias L0	0.0411±0.0150	0.0123±0.0083	0.0044±0.0037	0.0006±0.0005	M4_a+n L0	0.0953±0.0305	0.0749±0.0496	0.0443±0.0318	0.0048±0.0054
S4_alias L1	0.0843±0.0282	0.0504±0.0295	0.0242±0.0143	0.0009±0.0006	M4_a+n L1	0.1092±0.0394	0.0784±0.0480	0.0564±0.0399	0.0058±0.0061
S4_alias L2	0.1410±0.0411	0.0710±0.0376	0.0542±0.0215	0.0015±0.0009	M4_a+n L2	0.1705±0.0486	0.0969±0.0563	0.1113±0.0588	0.0115±0.0104
S4_alias L3	0.1763±0.0403	0.0847±0.0434	0.0926±0.0358	0.0020±0.0012	M4_a+n L3	0.1993±0.0483	0.1031±0.0585	0.1437±0.0676	0.0170±0.0160
S5_metal L0	0.0955±0.0229	0.1069±0.0498	0.1860±0.0515	0.0011±0.0007	M5_m+b+n L0	0.1745±0.0618	0.1051±0.0527	0.2013±0.0484	0.0040±0.0043
S5_metal L1	0.1022±0.0230	0.1106±0.0490	0.1952±0.0518	0.0012±0.0008	M5_m+b+n L1	0.2301±0.0848	0.1137±0.0577	0.2225±0.0525	0.0080±0.0073
S5_metal L2	0.1088±0.0233	0.1140±0.0488	0.1977±0.0520	0.0012±0.0008	M5_m+b+n L2	0.3264±0.0750	0.1272±0.0661	0.2666±0.0468	0.0168±0.0125
S5_metal L3	0.1140±0.0244	0.1158±0.0473	0.2000±0.0523	0.0013±0.0008	M5_m+b+n L3	0.3783±0.0579	0.1376±0.0679	0.3068±0.0448	0.0246±0.0158

Setting keywords: S1-S5 = single degradations; M1-M5 = mixtures; L0-L3 = severity levels.