Appendix

1. Severity distributions for independent corruptions

In the case of independent corruptions, two variants of the SCM were created which only differed by the distributions over corruption parameters. Each parameter contributes directly to corruption severity and in all cases, larger parameter values produce larger corruptions.

Table 1. Corruption parameter distributions for the independent corruptions SCM

		Distribution	
Corruption	Parameter	Uniform	Non-uniform
Gamma	γ	$(1, 3)$	(1)
Blur	σ	$(1, 11)$	(11)
Defocus	z, f_{stop}	$(1, 10)$, $(64, 128)$	(3), $(64, 128)$
Lens distort	distort, disperse	$(0, 0.1)$, $(0, 0.5)$	(0.3), (0.7)
Directional blur	distance	$(0, 0.1)$	(0.2)
Noise	scale	$(0, 0.25)$	(0.25)
Clouds	factor	$(0, 0.3)$	(0.3)
Glare	mix	$(-0.5, 0.5)$	$(0, 0.5)$

2. Causal model for dependent corruptions

The causal model used for Experiment 3 follows the structure in Figure 1. Each corruption parameter γ_i is a function of its parents and the exogenous noise ϵ_i (i.e., $\gamma_i = f(\gamma_{pa(i)}, \epsilon_i)$).

Figure 1. Causal model for non-IID image corruptions

The structural equations for this model are as follows:

Clouds:

$$\epsilon \sim \mathcal{U}(0, 1)$$
$$factor(\epsilon) = \begin{cases} 0 & \epsilon < 0.75 \\ x \sim \mathcal{HN}(0.3) & else \end{cases}$$

Blur:

$$\sigma(factor) = \begin{cases} 1 & factor > 0.2 \\ k \in \{1..9\},\ p(k) = 1/9 & else \end{cases}$$

Gamma:

$$\epsilon \sim \mathcal{U}(0, 1)$$
$$\gamma(k, \epsilon) = \begin{cases} 0.1 \cdot \epsilon + 1 & k \leq 3 \\ \epsilon + 1 & k > 3 \end{cases}$$

Lens distortion:

$$\epsilon \sim \mathcal{U}(0, 1)$$
$$distort(\gamma, \epsilon) = \begin{cases} 0.05 \cdot \epsilon & \gamma > 1.2 \\ 0.5 \cdot \epsilon & 1.0 < \gamma \leq 1.2 \end{cases}$$

Displacement/motion blur:

$$\epsilon_z, \epsilon_d \sim \mathcal{U}(0, 1)$$
$$zoom(distort, \epsilon_z) = 0.1 \cdot \epsilon_z$$
$$distance(distort, \epsilon_d) = 0.05 \cdot \epsilon_d$$

Defocus blur:

$$z(zoom, distance) = \begin{cases} x \in \{1..10\},\ p(x) = 1/10; & \\ \quad (zoom = 0) \vee (distance = 0) & \\ 1 \quad else & \end{cases}$$
$$f_stop(zoom, distance) = \begin{cases} x \in \{64..128\},\ p(x) = 1/64; & \\ \quad (zoom = 0) \vee (distance = 0) & \\ 128; \quad else & \end{cases}$$

Noise:

$$\epsilon \sim \mathcal{U}(0, 1)$$
$$\sigma_n(z, f_stop, \epsilon) = \begin{cases} 0.2 \cdot \epsilon & (z > 4) \vee (f_stop > 100) \\ 0.05 \cdot \epsilon & else \end{cases}$$

This structural causal model is meant to induce wide variability in the image corruptions in contrast to the IID model. The structural dependencies were specified to produce more visually complex corruptions. A comparison of Experiment 1-3 illustrates how OC model performance varies significantly with changes in the corruption generating process.

3. Object recovery vs. Severity by Corruption

The following figures compare the per-corruption performance of models as a function of normalized severity. In each case, normalized severity is determined by the range of the sampled corruption parameters. The top panel in each figure corresponds to the case where corruptions are sampled IID (Experiment 1) and the bottom panel corresponds to the non-IID case (Experiment 3).

The figures show that mIoU (object recovery) as a function of severity may vary significantly across algorithms. For instance, in Figure 6, both IODINE and SPAIR produce the highest clean performance, but mIoU drops rapidly with severity which is in stark contrast to the other methods analyzed.

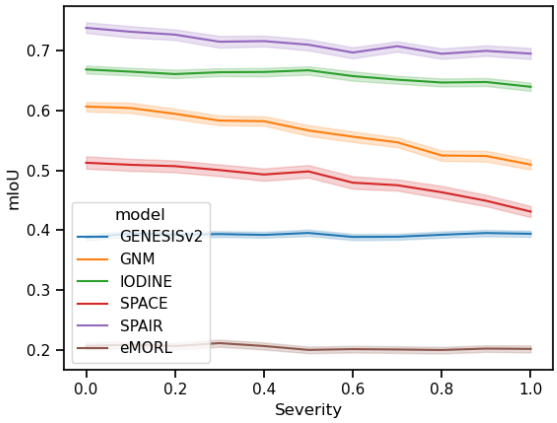

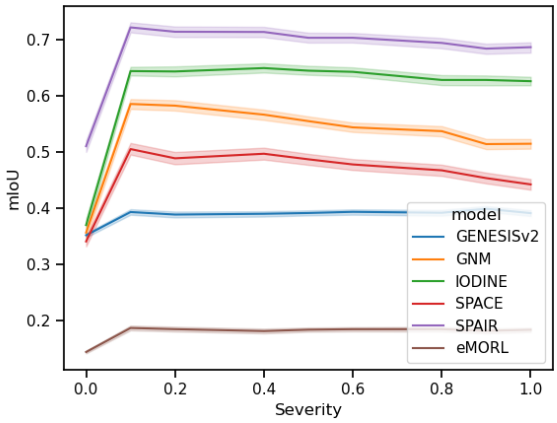

Figure 2. Blur corruption

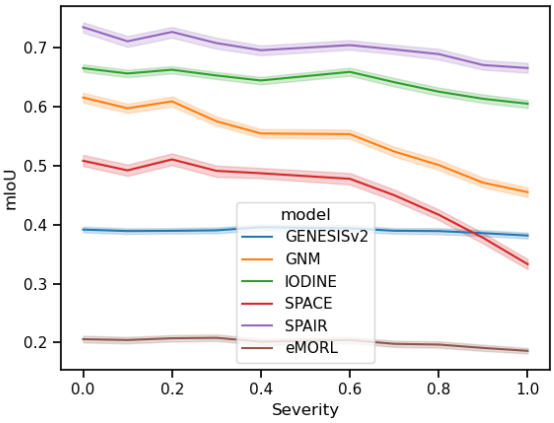

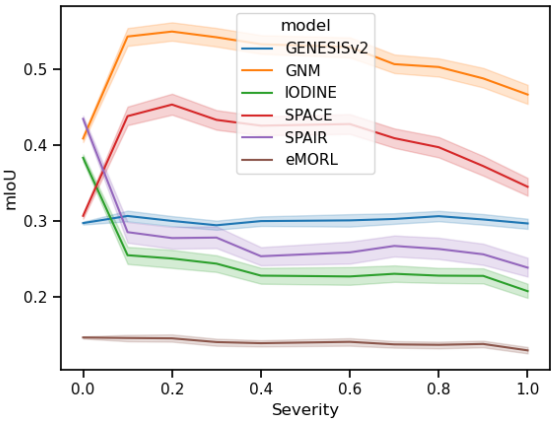

Figure 3. Defocus Blur

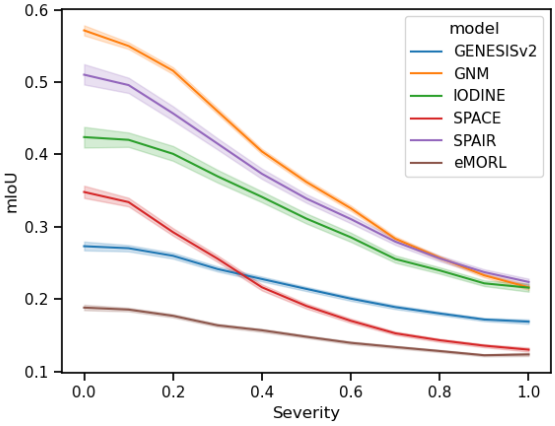

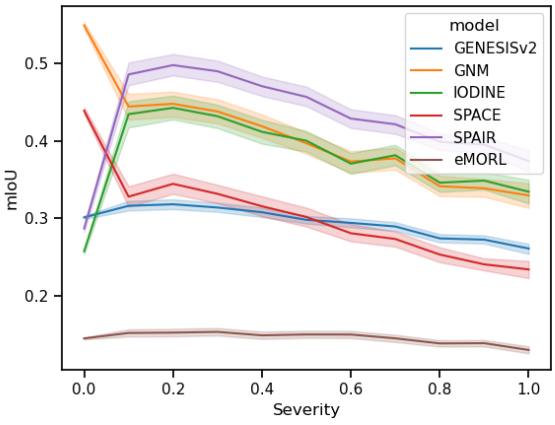

Figure 4. Displacement/motion blur

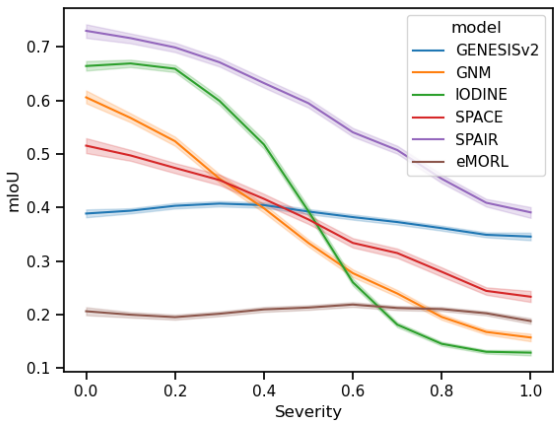

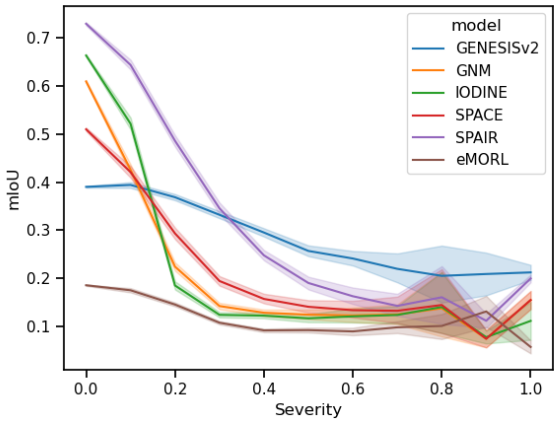

Figure 5. Clouds

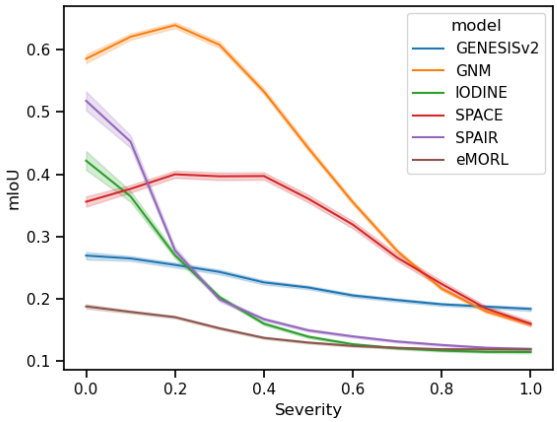

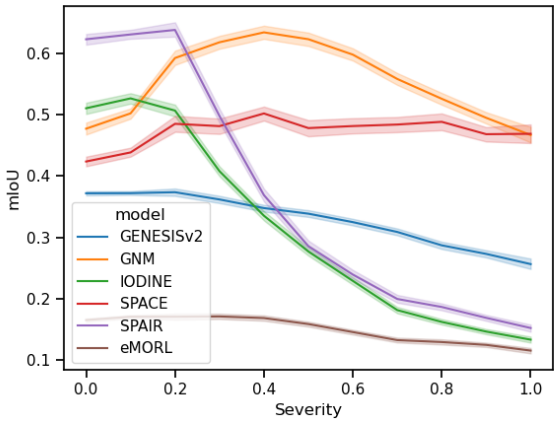

Figure 6. Gamma corruption

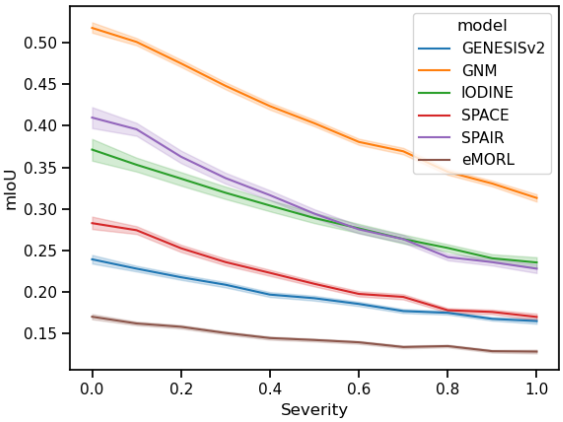

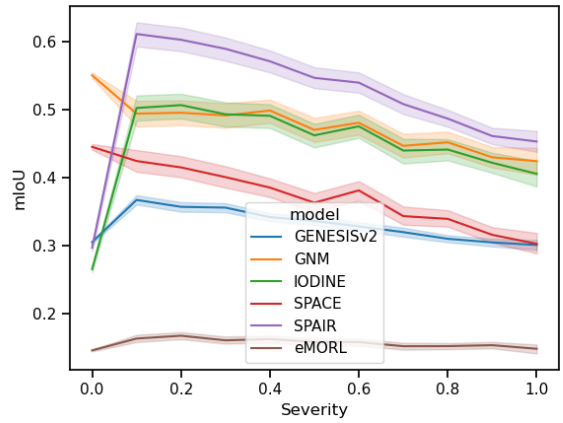

Figure 7. Lens Distortion

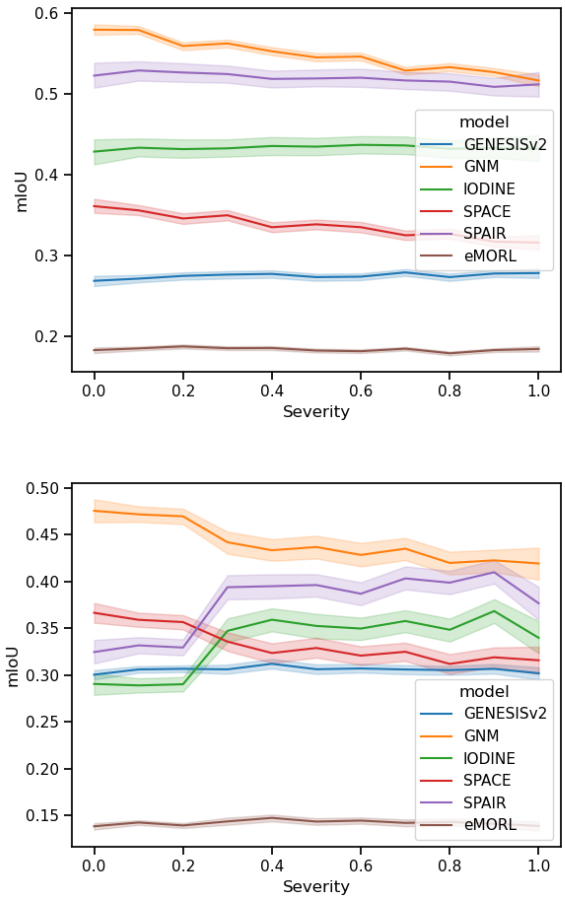

Figure 8. White noise corruption