Zoom in for further details

Table 1: Organization of reviewer comments.

Reviewer Questions	XxiF	j7y2	SjE4	ZG4p	fxDa
I. Novelty and contribution	W1				W2
II. Assuming causal diagram	W2, L1	W1			W1
III. Domain vs intervention	W1	Q1			
IV. Non-Markovianity	W1, W2				Q2
V. Image Experiments	L2	W1	W1, Q1	W1	

Figure 1: Three axes of identifiability in causal representation learning.

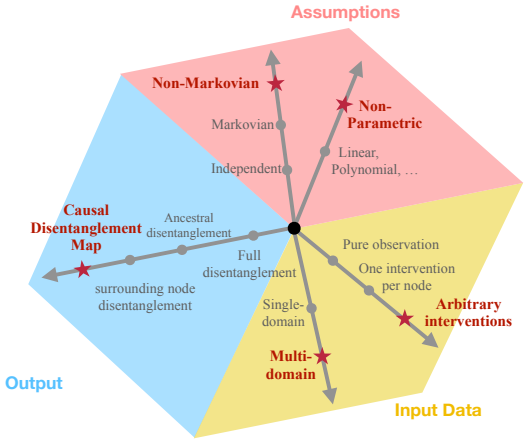

Figure 2: The underlying data generating process in ColoredMNIST and Bar experiments. **(A)** the LSD induced by the true ASCMs; **(B)** observations from domain Π_1; **(C)** observations from Π_2; **(D)** soft intervention on V_3 in Π_1, and **(E)** hard intervention on V_2 in Π_2. Gray arrows indicate invariant distributions w.r.t. (B). Red arrows indicate distribution changes w.r.t. (B). E.g., In (B), 9 is likely violet, which will influence the bar to be likely blue, whereas in (D), the bar is likely to be yellow.

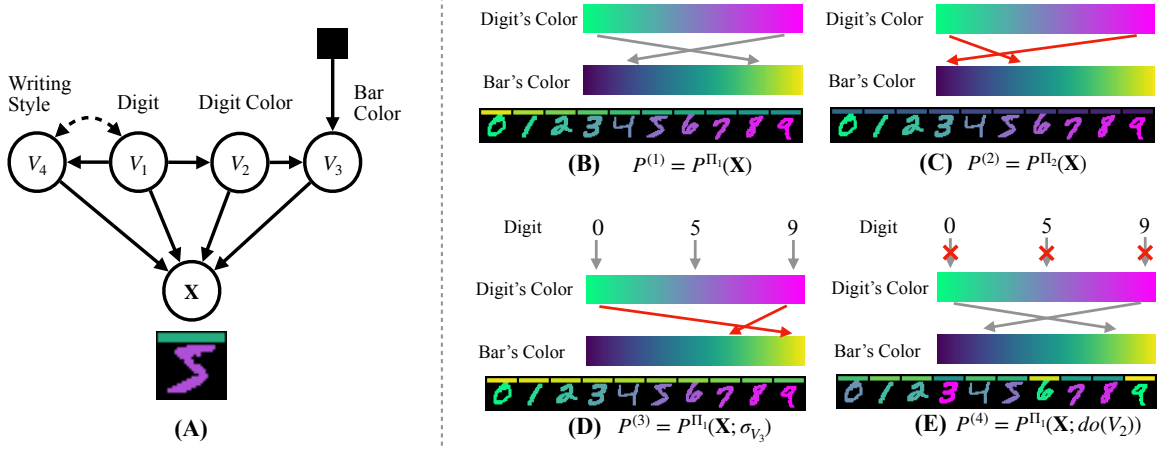

Figure 3: **(A)** The CDM output of CRID for Color-Digit and Color-Bar. Correlation of learned representation of Color-Digit **(B)** and Color-Bar **(C)** w.r.t. ground-truth latent variables. Pink denotes the output disentanglement predicted by the CDM. Lower correlation implies disentanglement and higher correlation implies entanglement. Since V_4 is not known, there is no ground-truth to assess against.

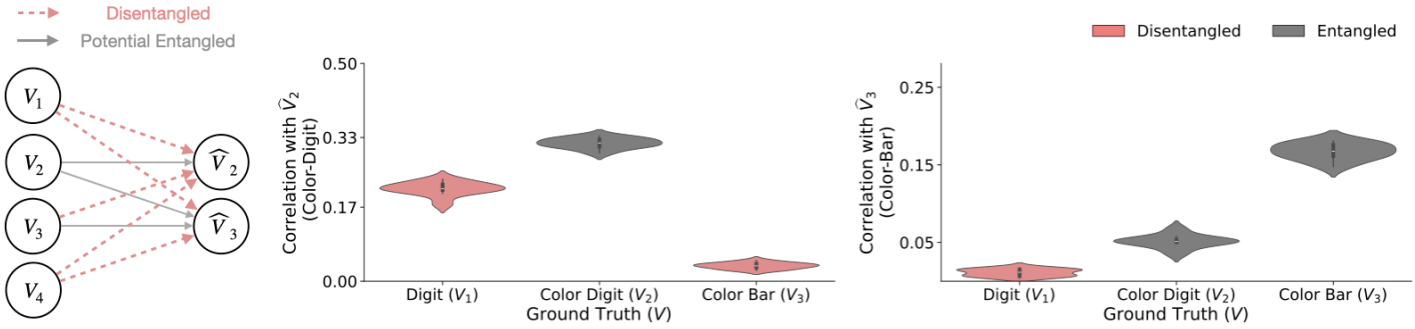

Figure 4: Image editing using learned representations of **(A)** digit color and **(B)** bar color. When editing digit color V_2 the digit number (V_1) and writing style (V_4) should not change while bar color V_3 can change because of the causal effect of V_2 on V_3). When editing the bar color, the the digit number (V_1) and writing style (V_4) should not change while the digit color V_2 can change because of the potential entanglement.