

The Force of Jensen-Shannon Divergence in Disentangling Diffusion Models

Eric Tillmann, Enis Simsar, Thomas Hofmann

ETH zürich

Problem:

Text-to-image models fail on multi-object prompts:

- Missing or omitted objects
- Attribute entanglement

Our Solution: JEDI

Jensen-Shannon Divergence for **D**isentangleme**n**t at **I**nference

- Refines generations on-the-fly → no retraining needed
- Lightweight adversarial steps → minimal compute overhead
- Plug-and-play with any model

Check it out at ericbill21.github.io/JEDI/!



StableDiffusion 3.5



"A dachshund and a corgi sitting [...]"

StableDiffusion 3.5
+ JEDI



"A horse, a bear and a moose in [...]"

