Conditional GANs / pix2pix

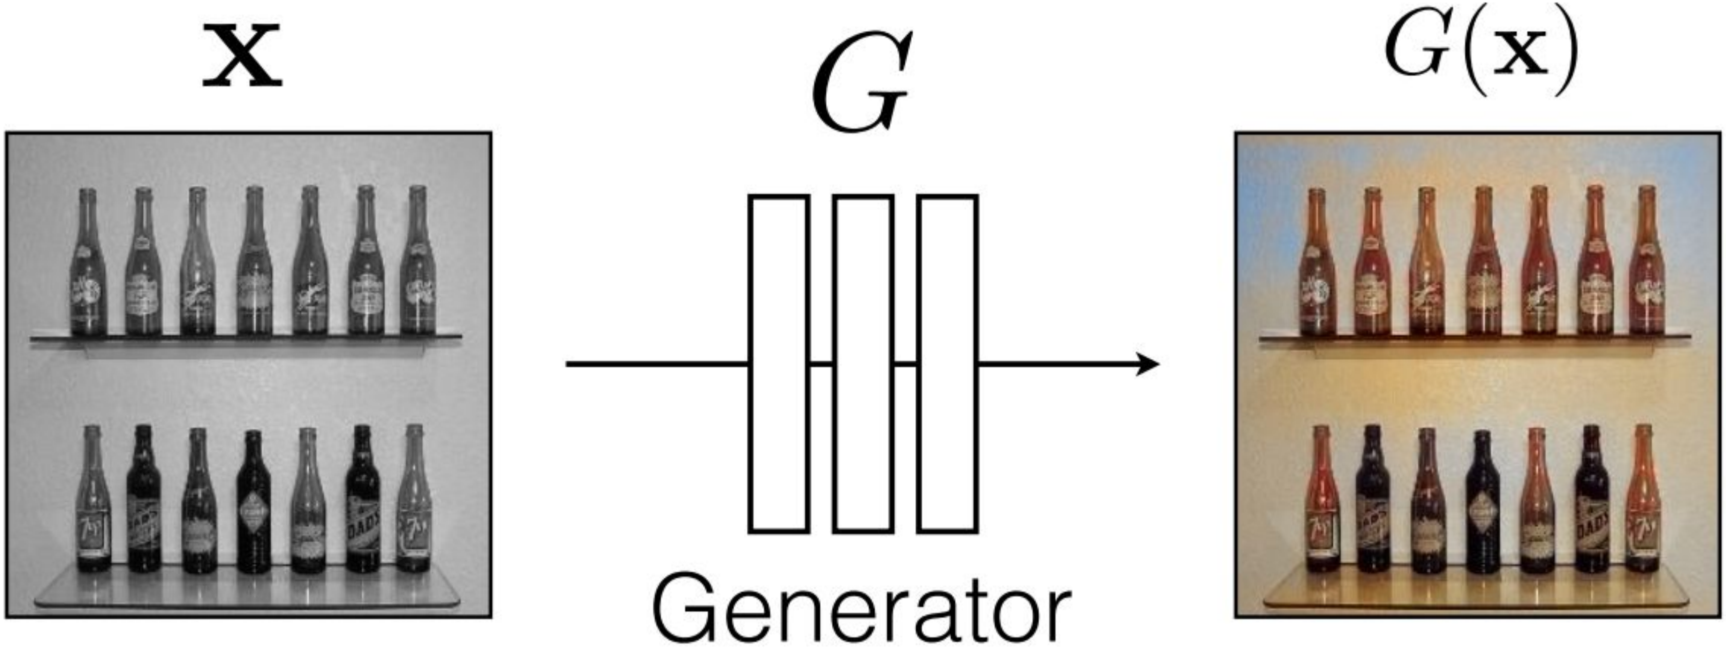

One-to-One Correspondence in the Domain is required: Colorization, Sketch to Image, etc

Slide: Phillip Isola

Conditional GANs / pix2pix

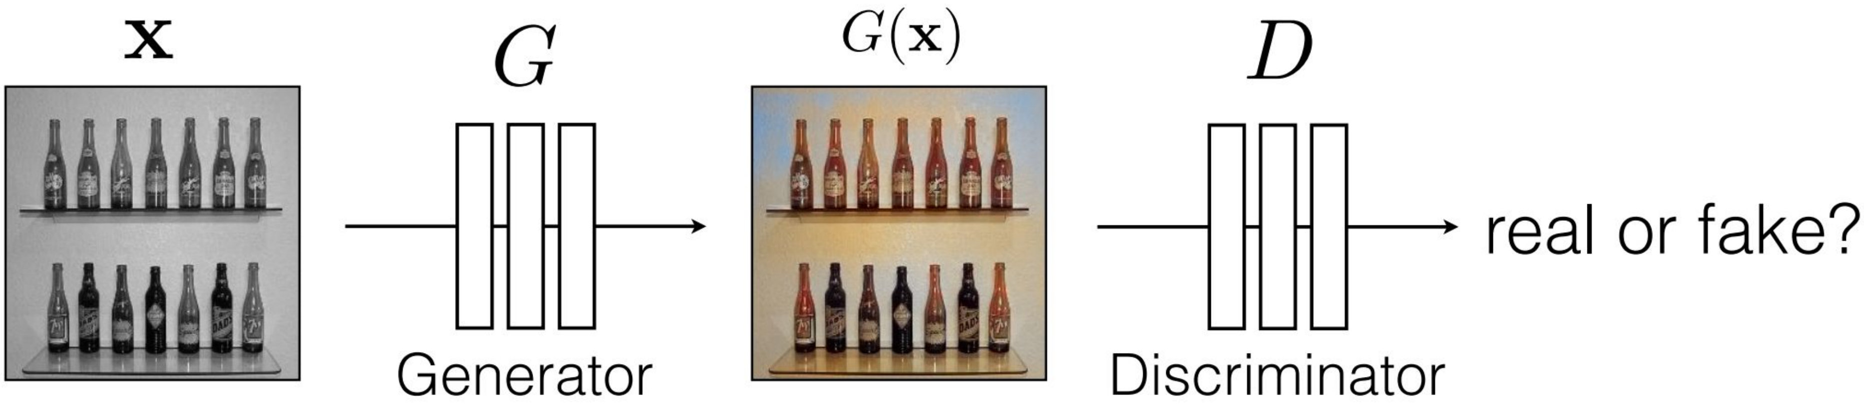

Slide: Phillip Isola

Conditional GANs / pix2pix

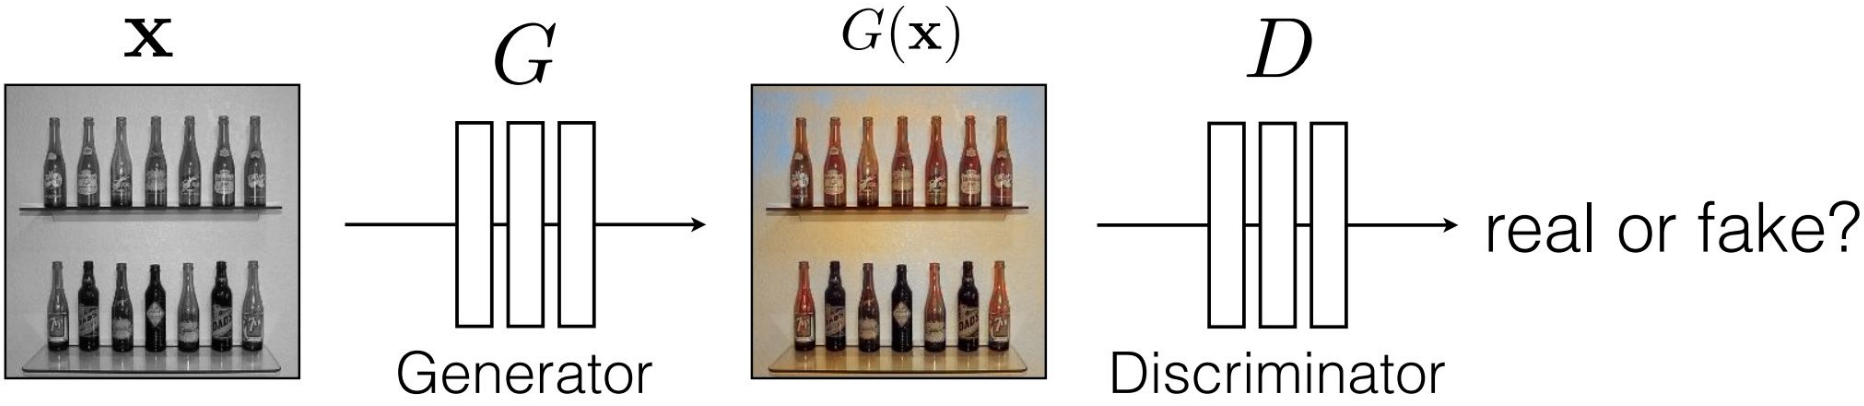

G tries to synthesize fake images that fool **D**

D tries to identify the fakes

Slide: Phillip Isola

Conditional GANs / pix2pix

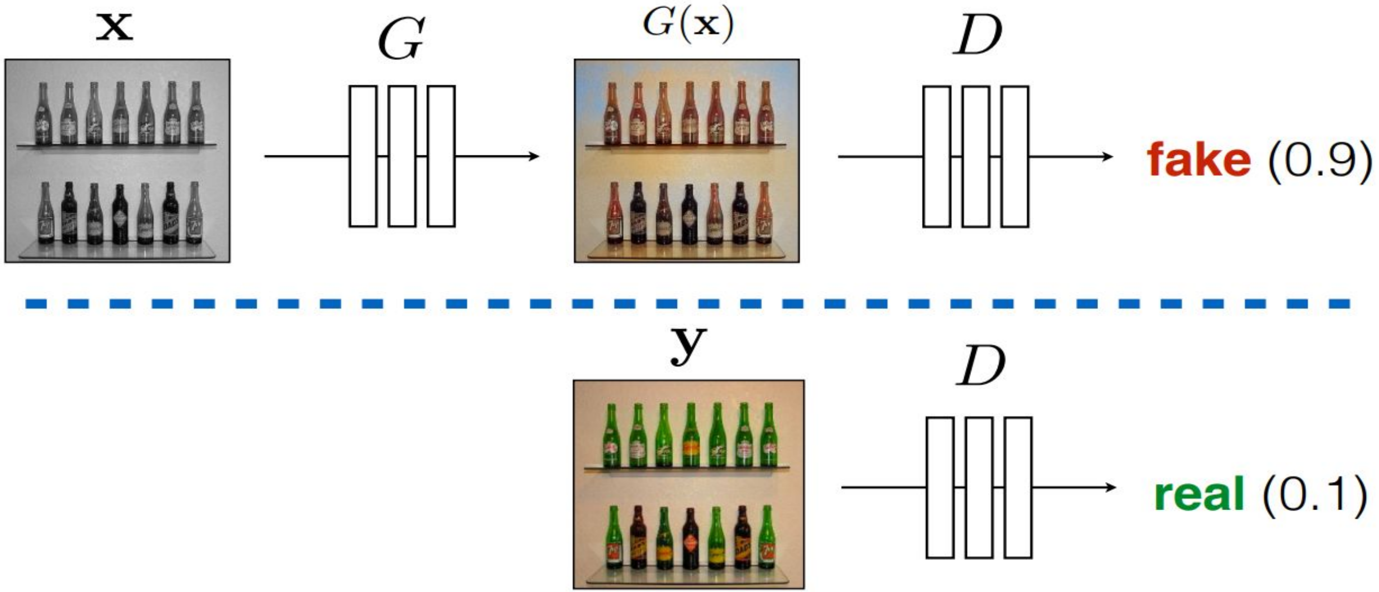

$$\arg\max_{D} \; \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \boxed{\log D(G(\mathbf{x}))} \ + \ \boxed{\log(1 - D(\mathbf{y}))}\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

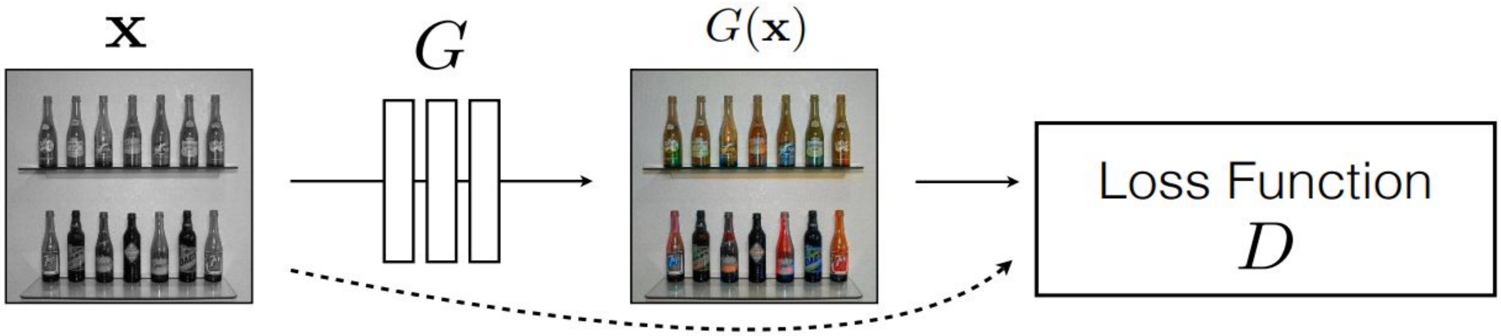

G's perspective: **D** is a loss function.

Rather than being hand-designed, it is *learned*.

Slide: Phillip Isola

Conditional GANs / pix2pix

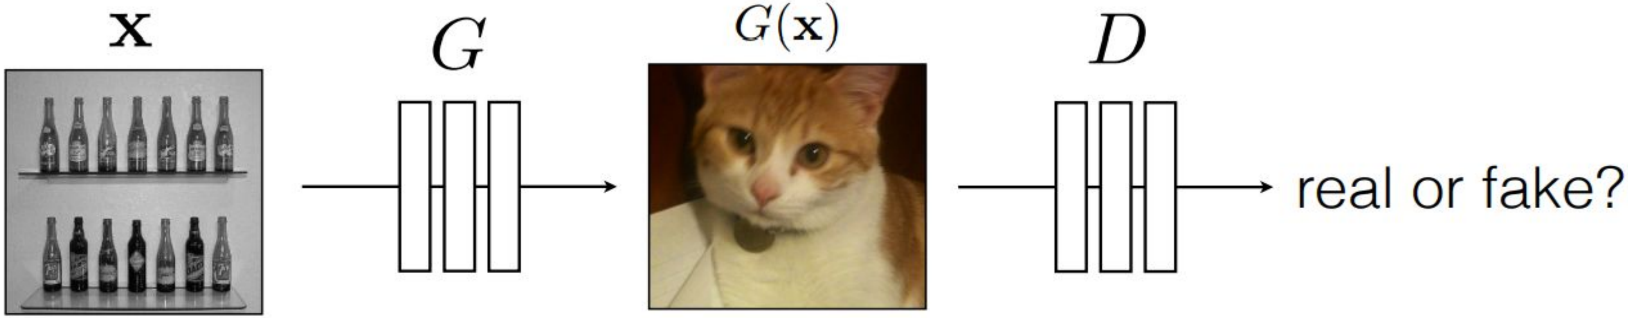

$$\arg\min_G \max_D \ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(G(\mathbf{x})) \ + \ \log(1 - D(\mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

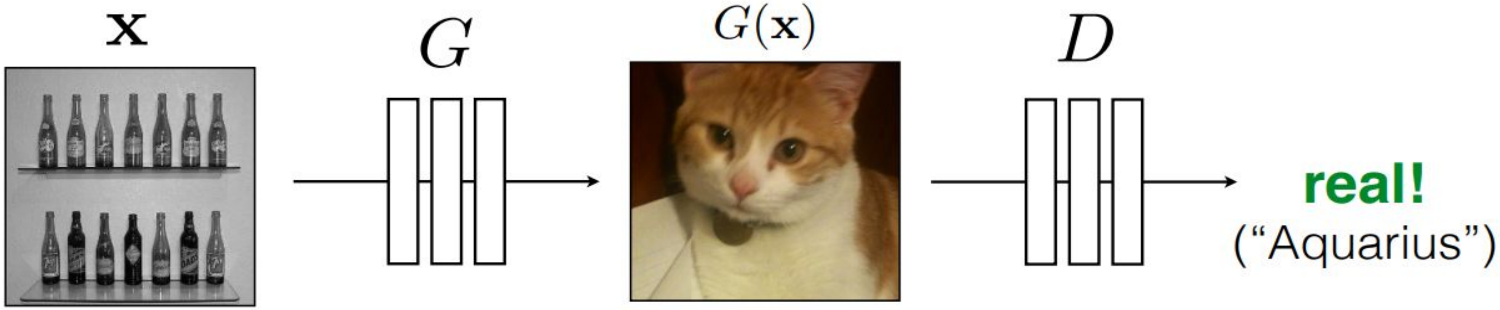

$$\arg\min_G \max_D \; \mathbb{E}_{\mathbf{x},\mathbf{y}}[\; \log D(G(\mathbf{x})) \;+\; \log(1 - D(\mathbf{y})) \;]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

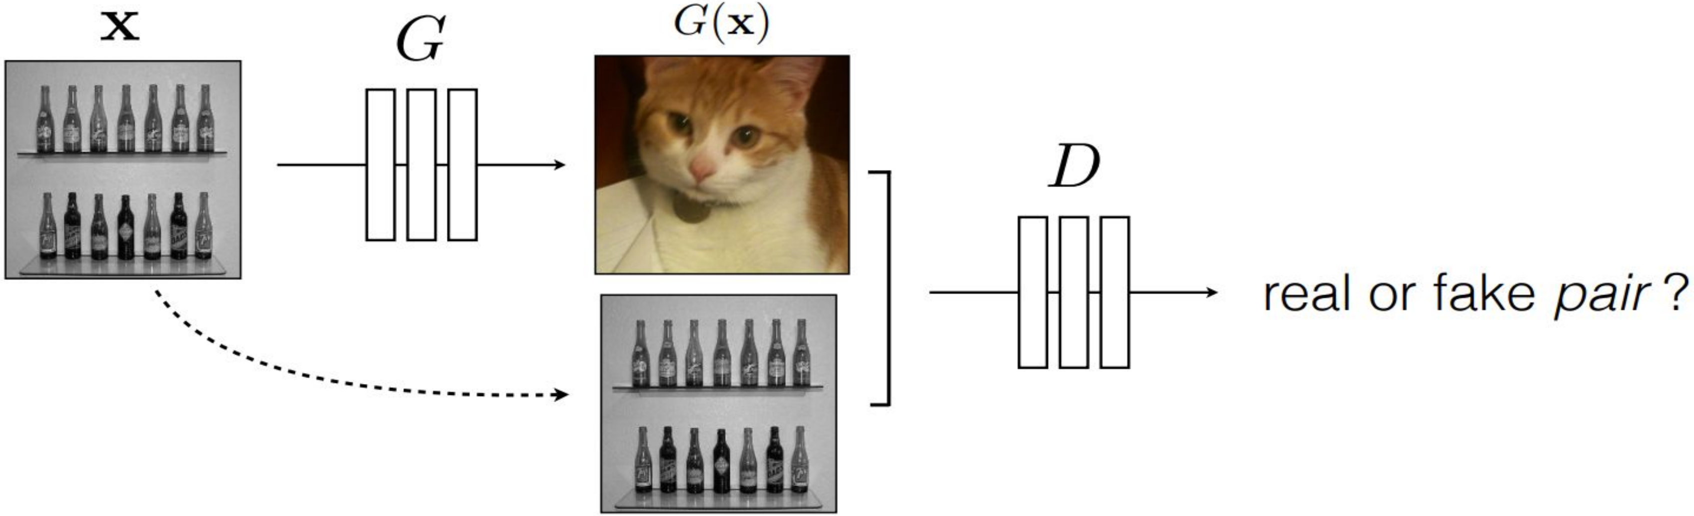

$$\arg\min_G \max_D \ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(G(\mathbf{x})) \ + \ \log(1 - D(\mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

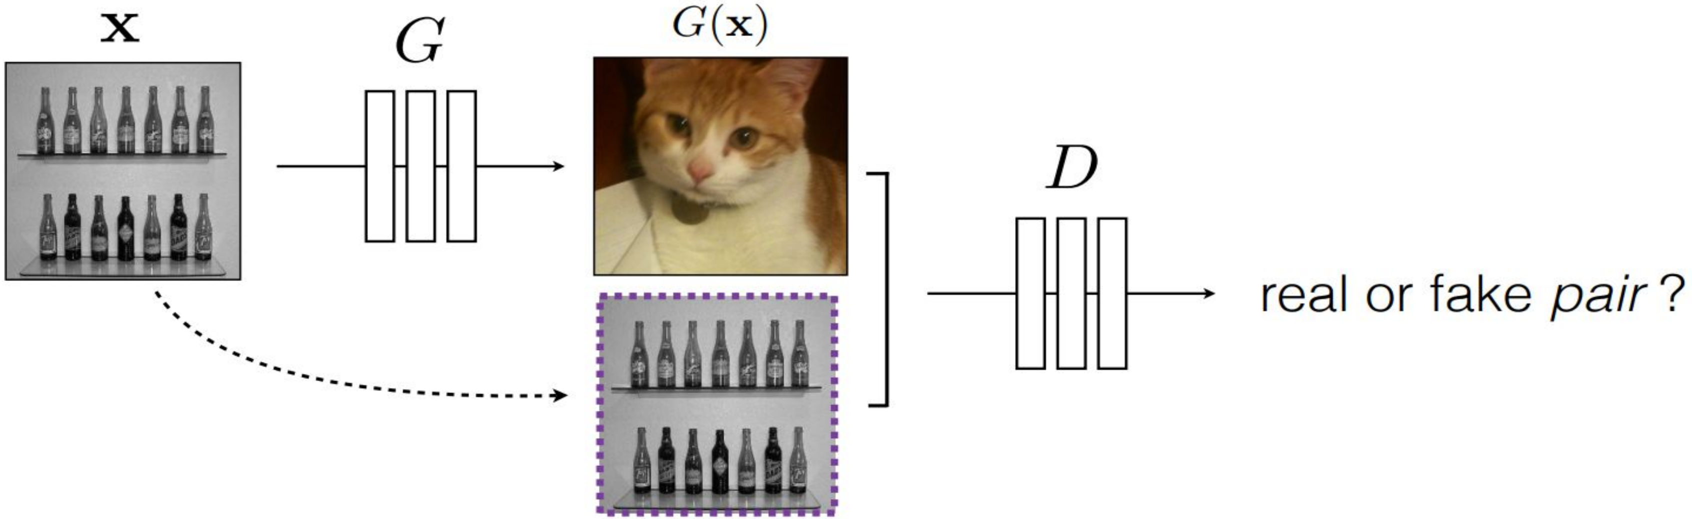

$$\arg\min_G \max_D \ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(\boxed{\mathbf{x}}, G(\mathbf{x})) + \log(1 - D(\boxed{\mathbf{x}}, \mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

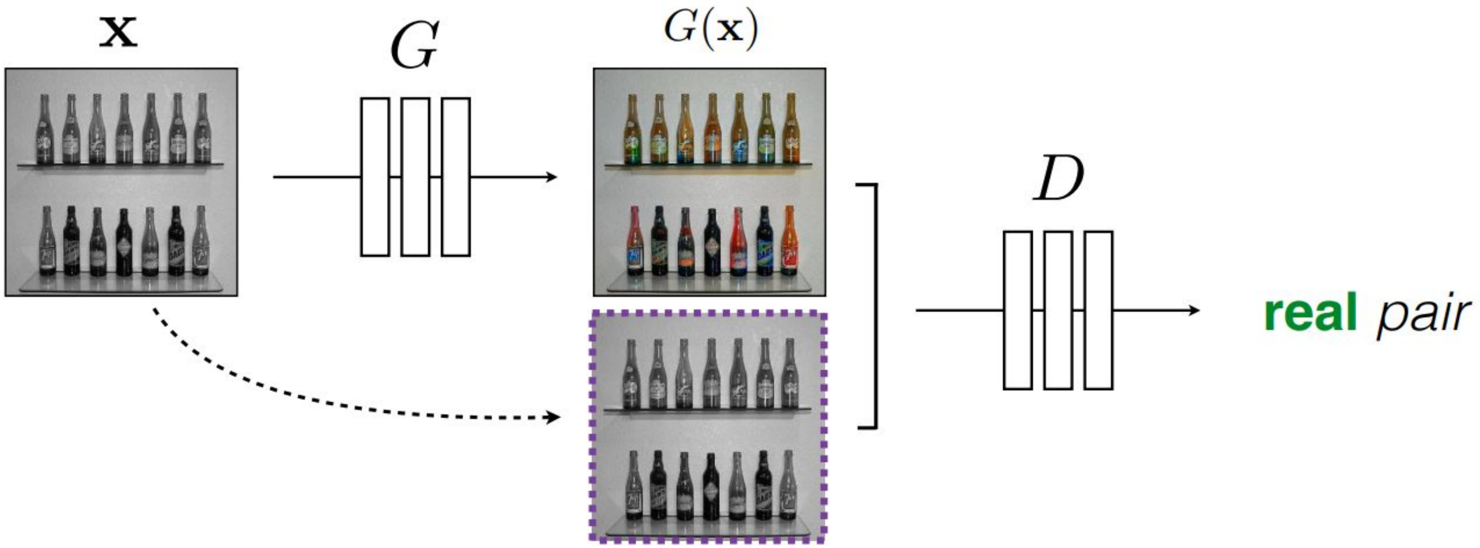

$$\arg\min_G \max_D \ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(\boxed{\mathbf{x}}, G(\mathbf{x})) + \log(1 - D(\boxed{\mathbf{x}}, \mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

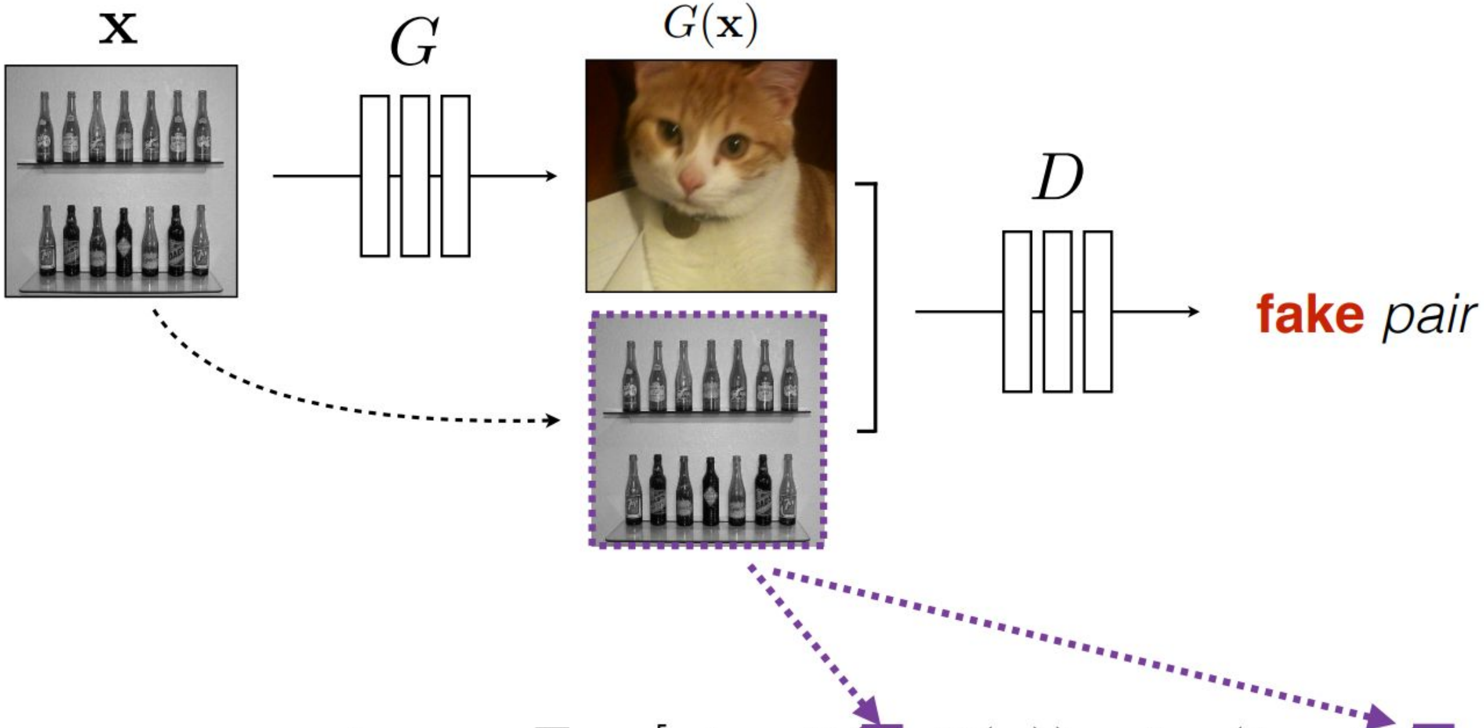

$$\arg\min_{G}\max_{D}\ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(\boxed{\mathbf{x}}, G(\mathbf{x})) + \log(1 - D(\boxed{\mathbf{x}}, \mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

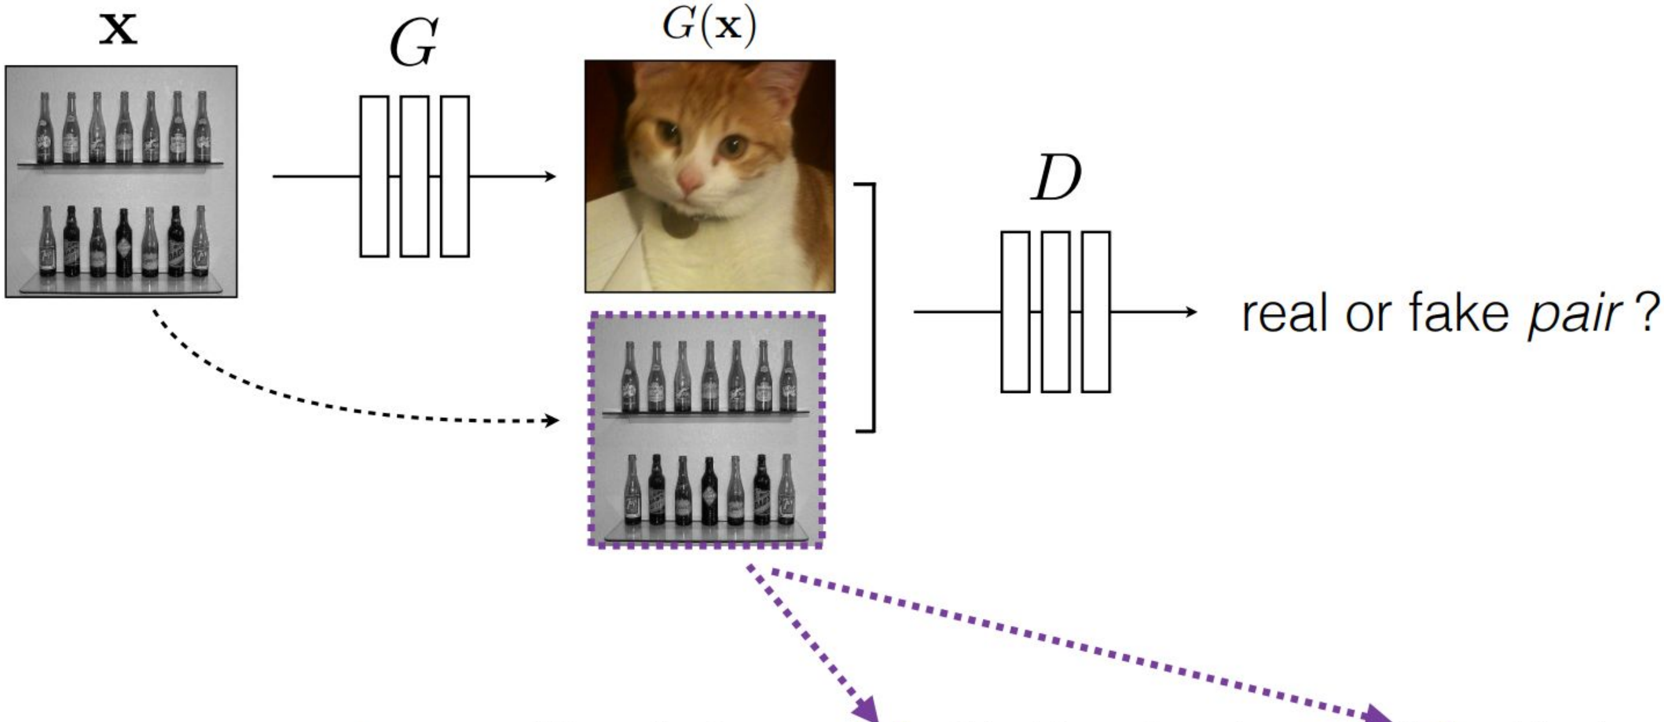

$$\arg\min_{G}\max_{D}\ \mathbb{E}_{\mathbf{x},\mathbf{y}}[\ \log D(\boxed{\mathbf{x}}, G(\mathbf{x})) + \log(1 - D(\boxed{\mathbf{x}}, \mathbf{y}))\ ]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

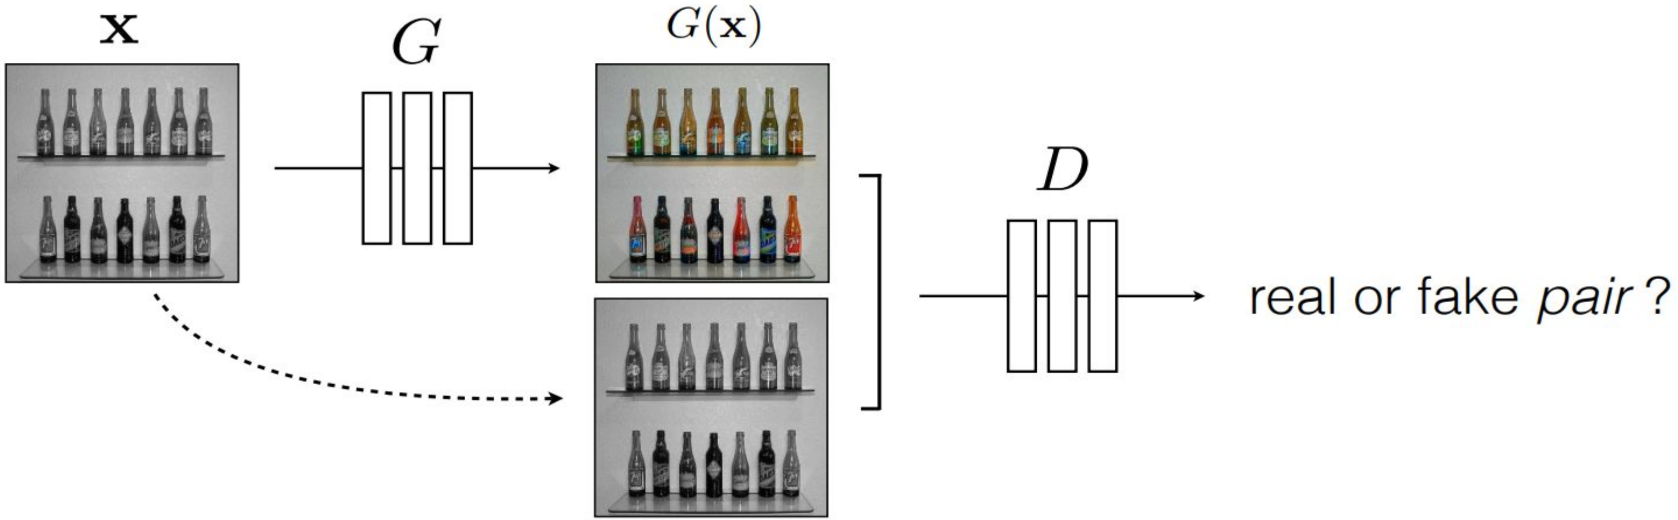

$$\arg\min_G \max_D \; \mathbb{E}_{\mathbf{x},\mathbf{y}}[\;\log D(\mathbf{x}, G(\mathbf{x})) + \log(1 - D(\mathbf{x}, \mathbf{y}))\;]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

$$G^* = \arg\min_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G)$$

Reconstruction loss is added

$$\mathcal{L}_{L1}(G) = \mathbb{E}_{x,y,z}[\|y - G(x, z)\|_1]$$

Slide: Phillip Isola

Conditional GANs / pix2pix

Shrinking the capacity: Patch Discriminator

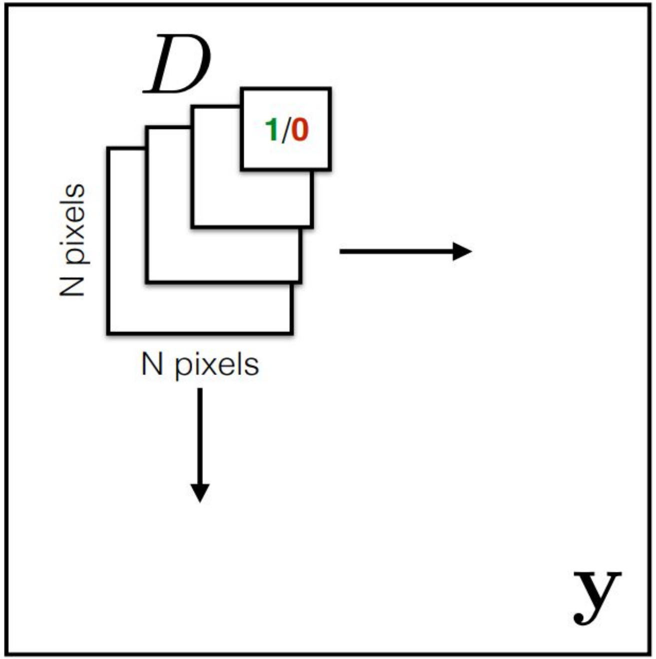

Rather than penalizing if output *image* looks fake, penalize if each overlapping *patch* in output looks fake

[Li & Wand 2016]
[Shrivastava et al. 2017]
[Isola et al. 2017]

Slide: Phillip Isola

Conditional GANs / pix2pix

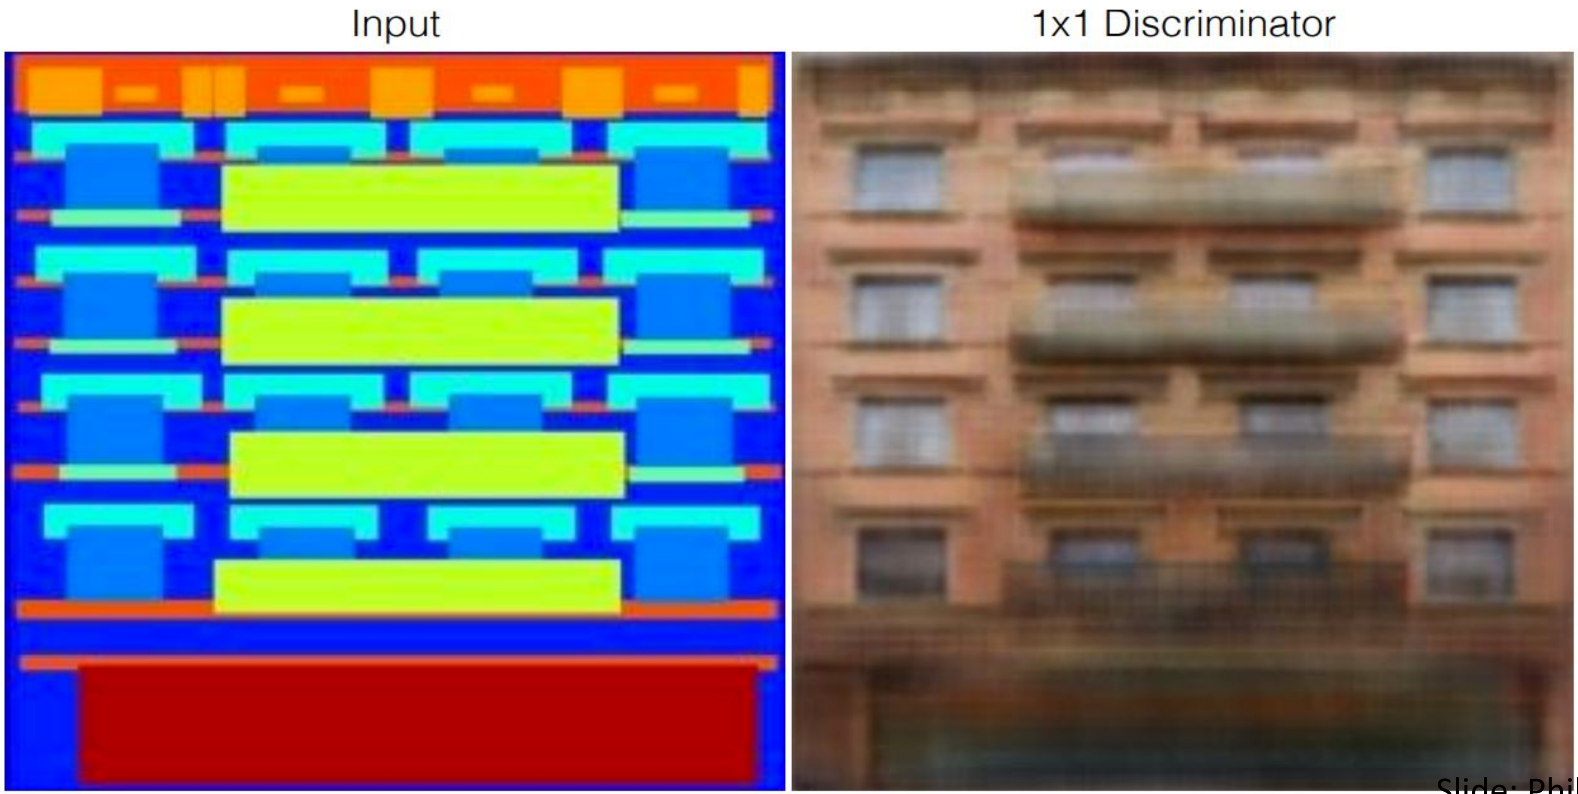

Slide: Phillip Isola

Conditional GANs / pix2pix

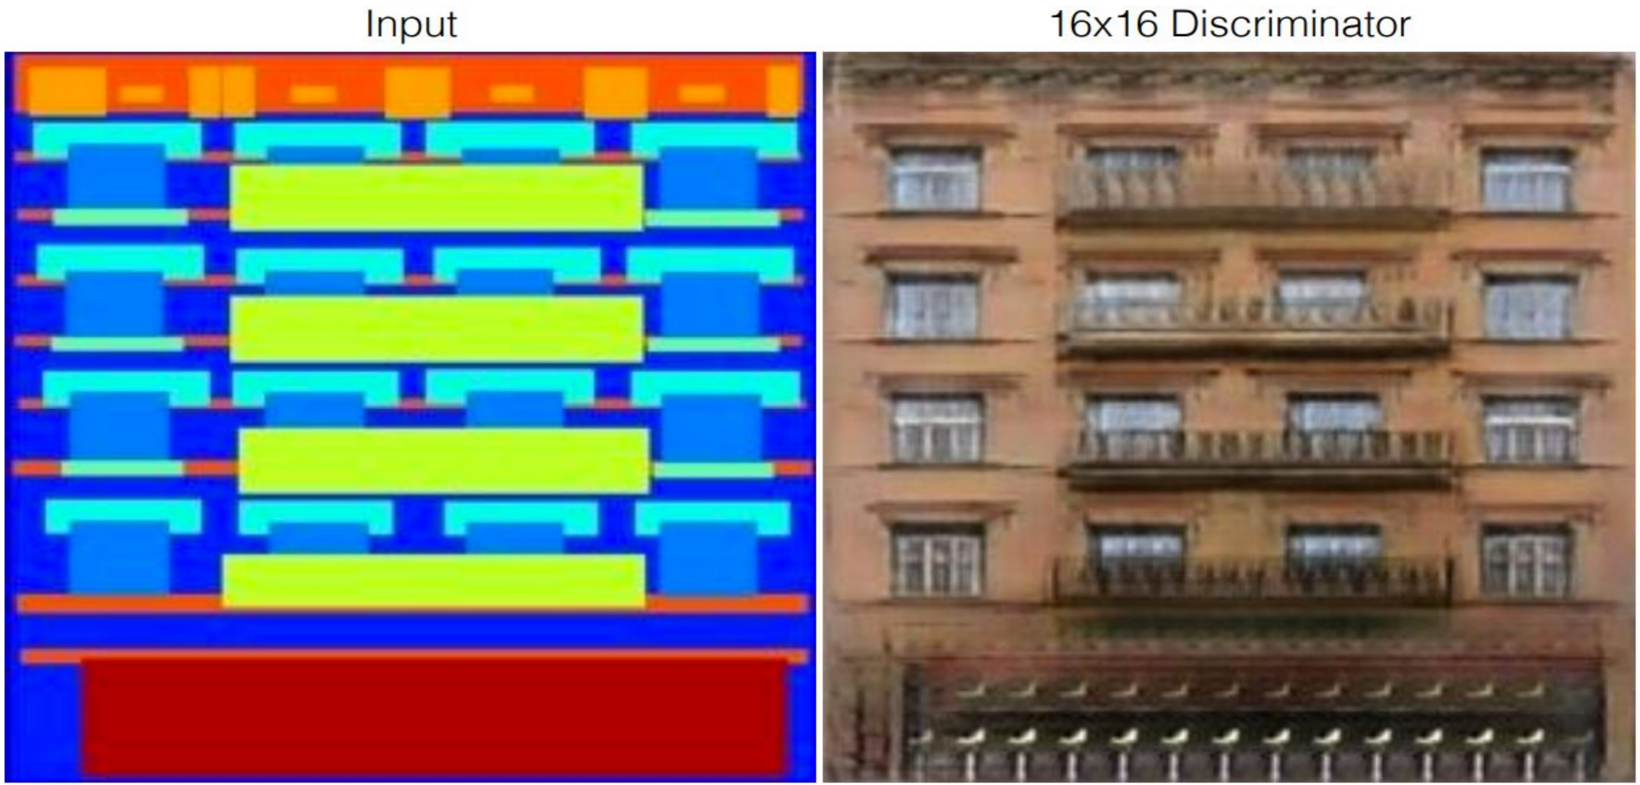

Slide: Phillip Isola

Conditional GANs / pix2pix

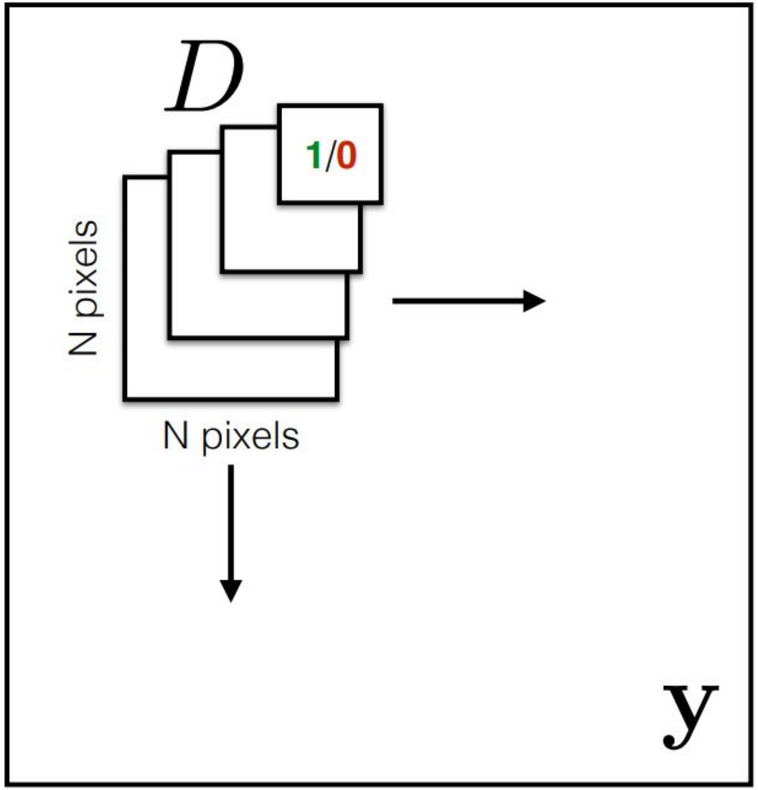

Rather than penalizing if output *image* looks fake, penalize if each overlapping *patch* in output looks fake

- Faster, fewer parameters
- More supervised observations
- Applies to arbitrarily large images

Slide: Phillip Isola

Conditional GANs / pix2pix

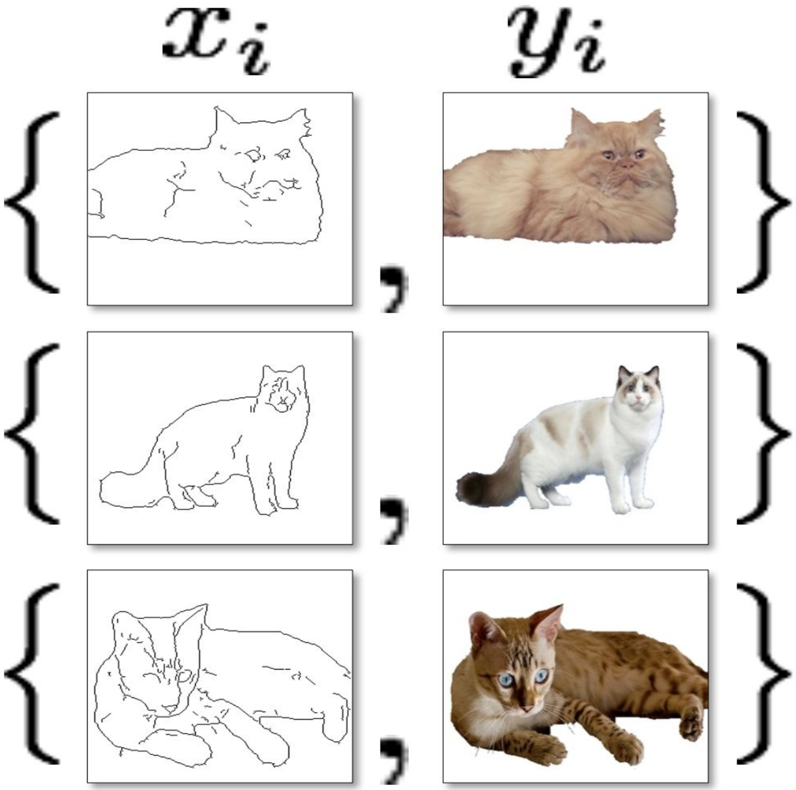

Slide: Phillip Isola

Conditional GANs / pix2pix

#edges2cats [Chris Hesse]

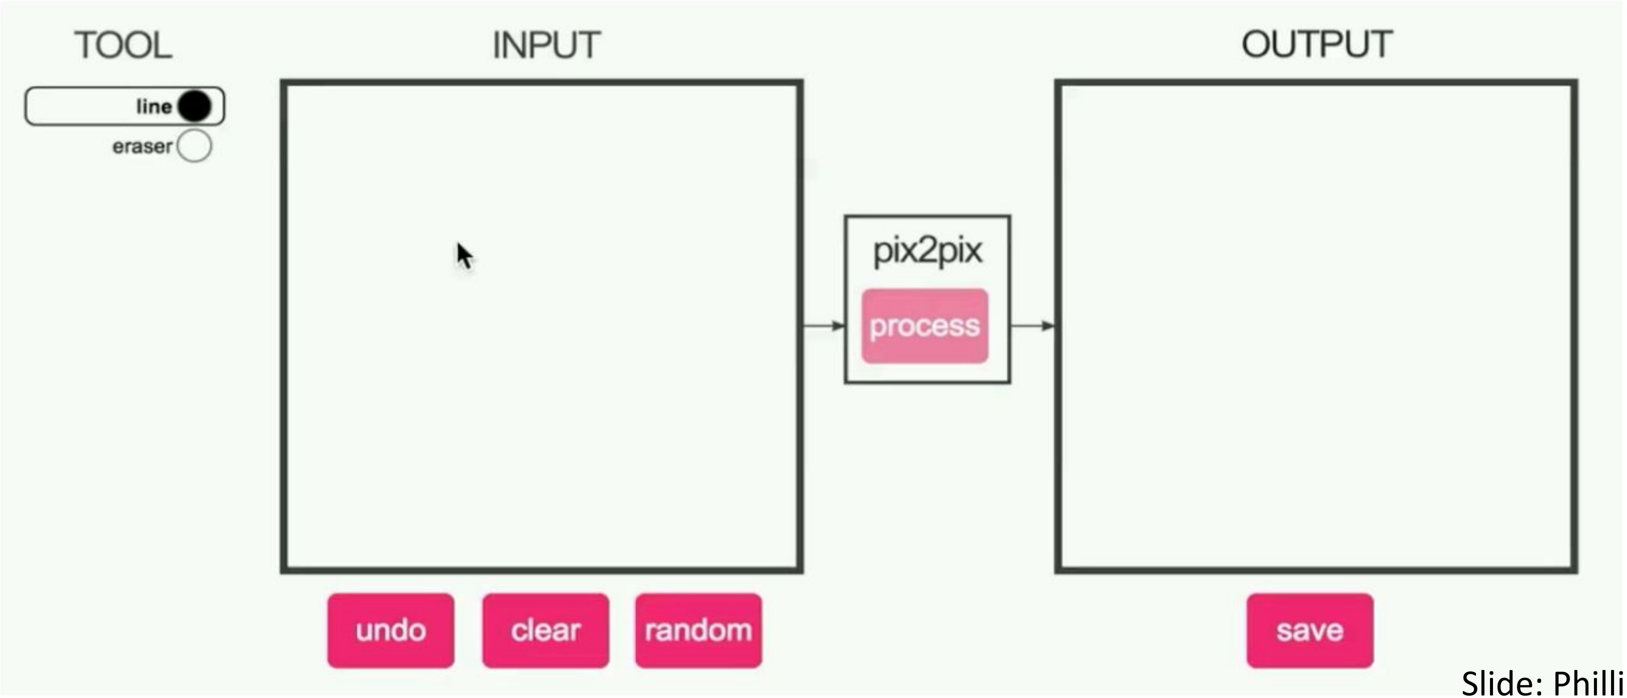

Slide: Phillip Isola

Conditional GANs / pix2pix

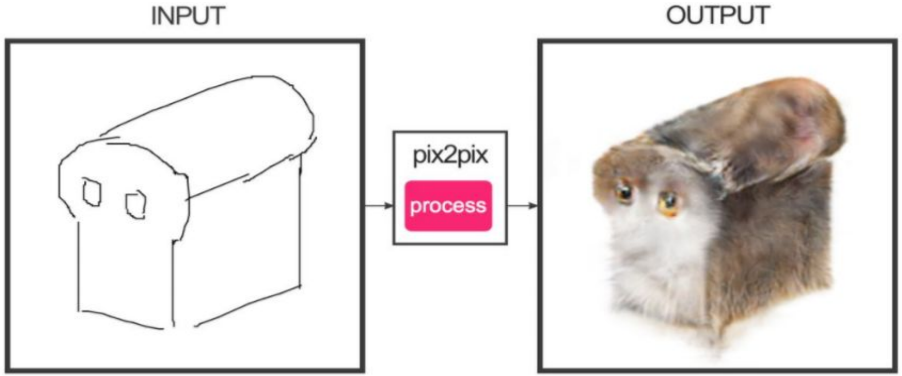

Ivy Tasi @ivymyt

Vitaly Vidmirov @vvid

Slide: Phillip Isola

Conditional GANs / pix2pix

BW → Color

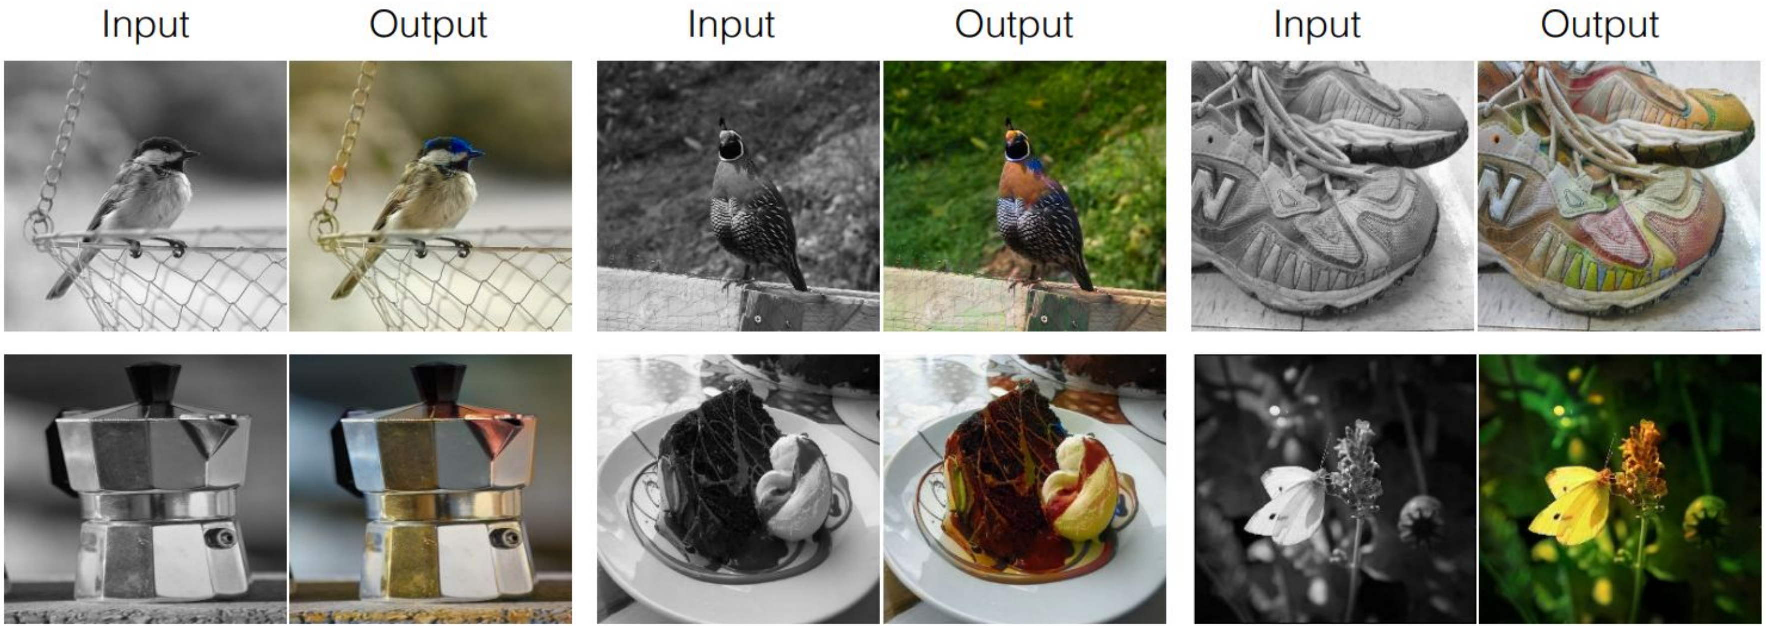

Slide: Phillip Isola