

Introduction to Deep Learning


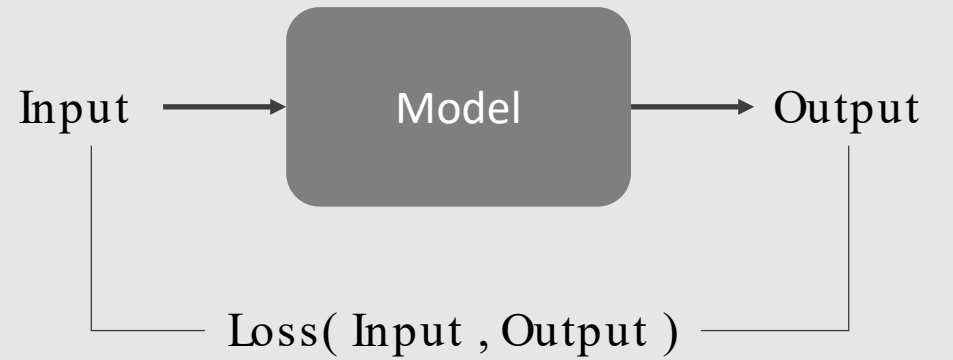
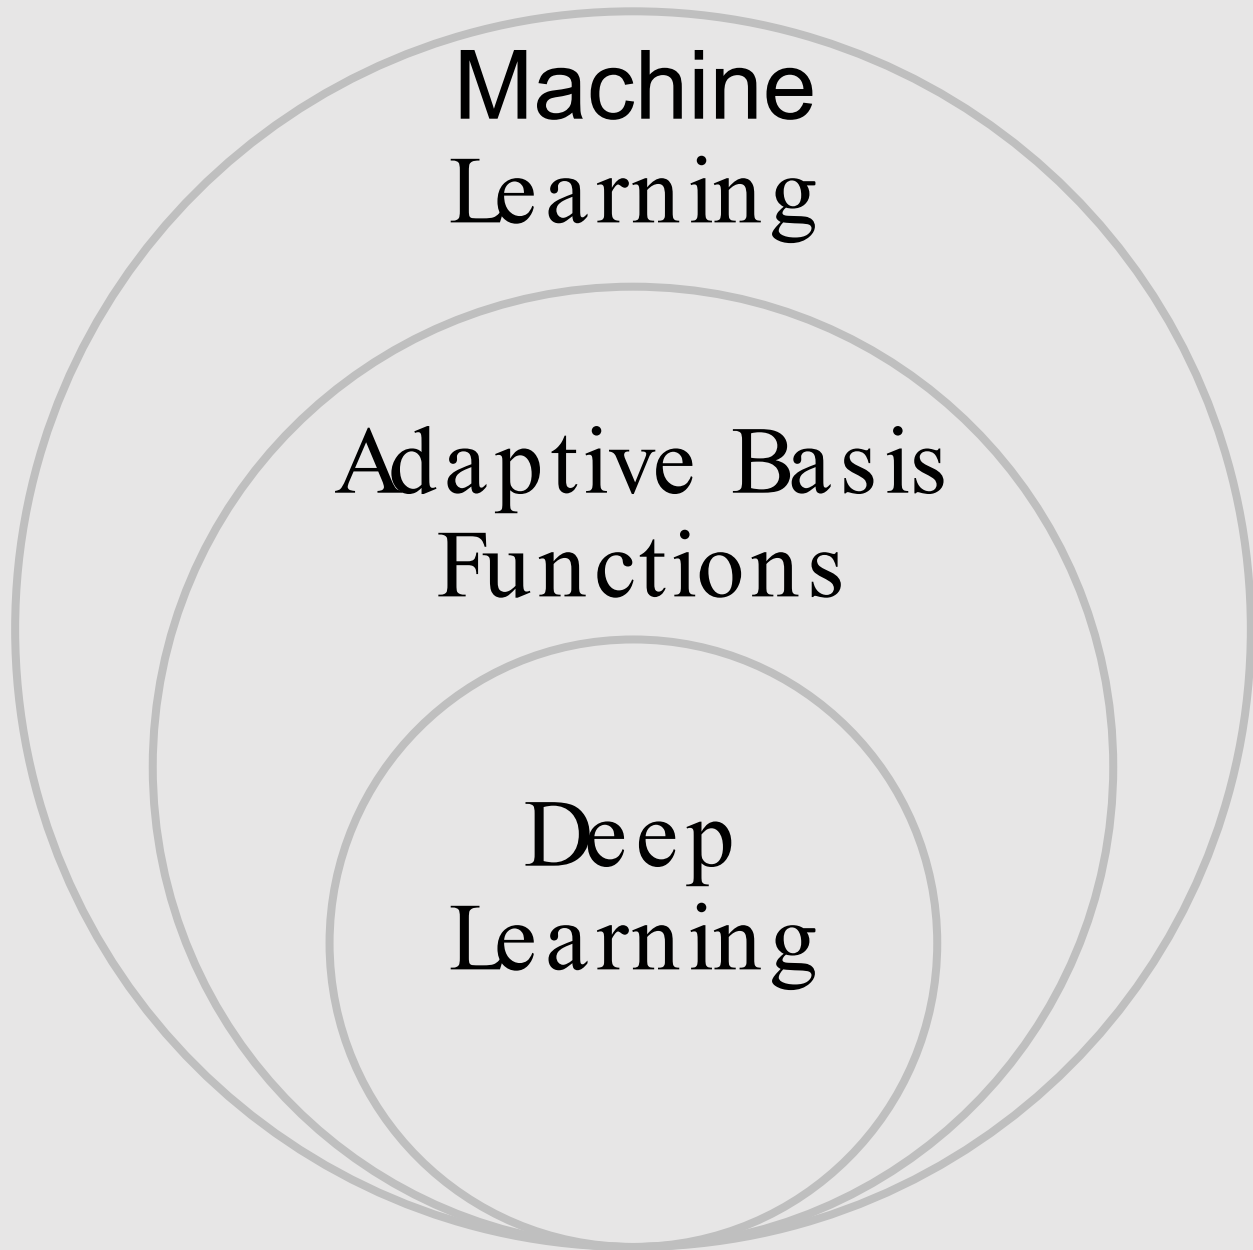
Dr. Mark Fuge

Univ. of Maryland, College Park

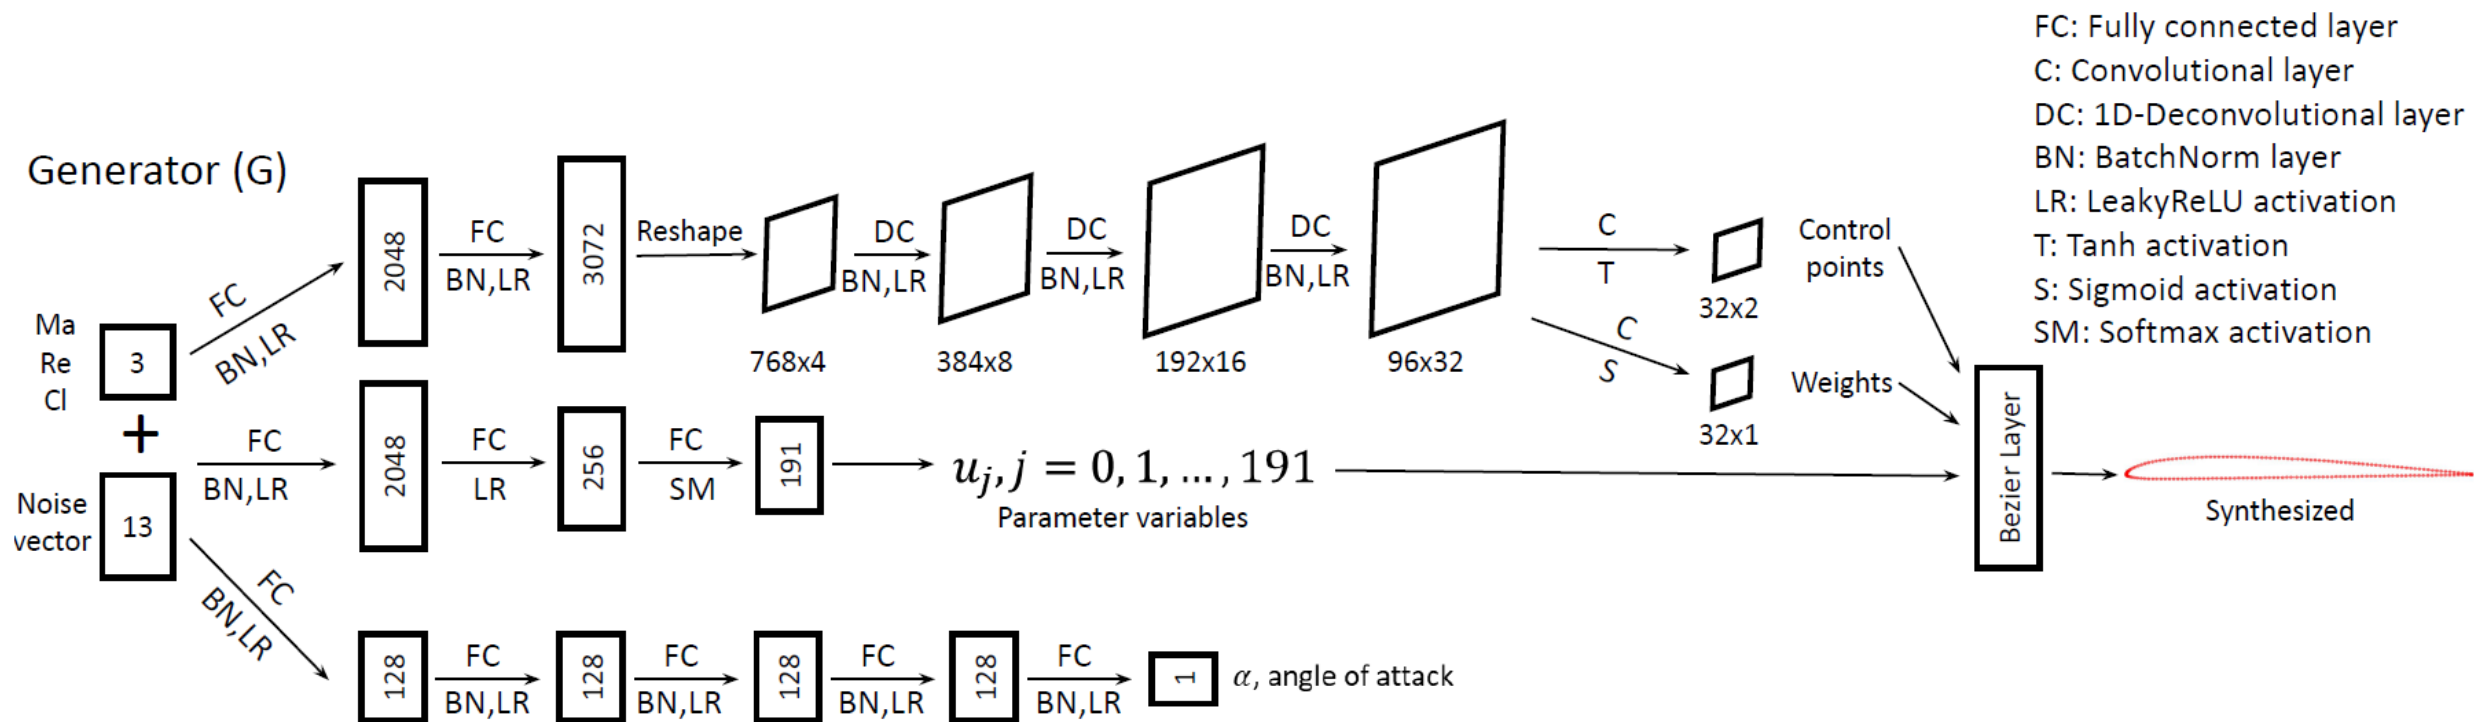
(301) 405-2558

fuge@umd.edu

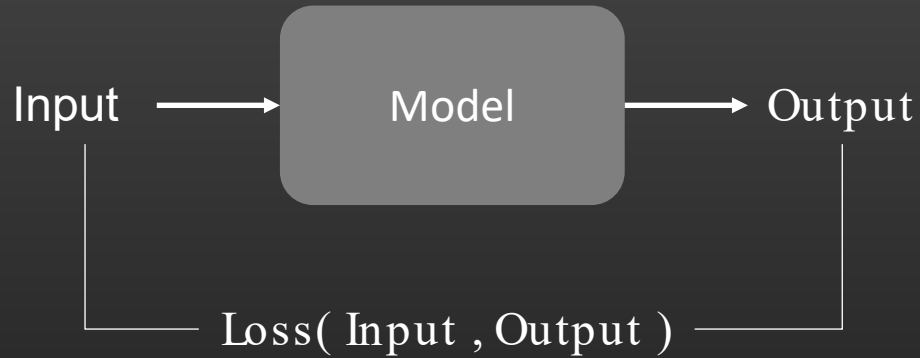
ideal.umd.edu



Typical unreadable Deep Learning Slide from my research group for a recent DoE Technical Review



Let's build a Deep Learning model to predict airfoil lift



What should the input be?
(This will be our *basis function*)

$$\text{Lift} = \text{Model}(\text{Input})$$
$$y = f(x)$$

How do you mathematically *represent* an airfoil?



How do you mathematically *represent* an airfoil?

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{199} \\ w_{200} \end{pmatrix}^T \begin{pmatrix} x_1 \\ y_1 \\ \vdots \\ x_{100} \\ y_{100} \end{pmatrix}$$

$\{x_1, y_1\}$




$$\begin{aligned} \text{Loss} &= (\text{Lift}_{\text{predicted}} - \text{Lift}_{\text{actual}})^2 \\ &= (\mathbf{w}^T \mathbf{x} - \text{Lift}_{\text{actual}})^2 \end{aligned}$$

Find w where:

$$\partial \text{Loss} / \partial w = 0$$

How do you mathematically *represent* an airfoil?

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \\ w_{11} \\ w_{12} \end{pmatrix}^T \begin{pmatrix} cx_1 \\ cy_1 \\ \\ t_3 \\ t_4 \end{pmatrix}$$



Only thing we changed was the *basis*.

But the basis was fixed/ static.

What if we *adapted* or *learned* the basis?

How do you *adapt* a basis?

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{199} \\ w_{200} \end{pmatrix}^T \begin{pmatrix} x_1 \\ y_1 \\ \vdots \\ x_{100} \\ y_{100} \end{pmatrix}$$



$$\begin{aligned} \text{Loss} &= (\text{Lift}_{\text{predicted}} - \text{Lift}_{\text{actual}})^2 \\ &= (\mathbf{w}^T \mathbf{x} - \text{Lift}_{\text{actual}})^2 \end{aligned}$$

How do you *adapt* a basis?

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{199} \\ w_{200} \end{pmatrix}^T \begin{pmatrix} g(x_1) \\ g(y_1) \\ \vdots \\ g(x_{100}) \\ g(y_{100}) \end{pmatrix}$$



$$\begin{aligned} \text{Loss} &= (\text{Lift}_{\text{predicted}} - \text{Lift}_{\text{actual}})^2 \\ &= (\mathbf{w}^T \mathbf{g}(\mathbf{x}) - \text{Lift}_{\text{actual}})^2 \end{aligned}$$


How do you *adapt* a basis?

What is g ?

Another model!

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{199} \\ w_{200} \end{pmatrix}^T \begin{pmatrix} g(x_1) \\ g(y_1) \\ \vdots \\ g(x_{100}) \\ g(y_{100}) \end{pmatrix}$$

$\{x_1, y_1\}$



$$\begin{aligned} \text{Loss} &= (\text{Lift}_{\text{predicted}} - \text{Lift}_{\text{actual}})^2 \\ &= (w^T g(x) - \text{Lift}_{\text{actual}})^2 \\ &= (w_1^T \{w_2^T(x)\} - \text{Lift}_{\text{actual}})^2 \end{aligned}$$


How do you *adapt* a basis?

What is g ?

Another model!

$$\text{Lift} = \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{199} \\ w_{200} \end{pmatrix}^T \begin{pmatrix} g(x_1) \\ g(y_1) \\ \vdots \\ g(x_{100}) \\ g(y_{100}) \end{pmatrix}$$

$\{x_1, y_1\}$




$$\begin{aligned} \text{Loss} &= (\text{Lift}_{\text{predicted}} - \text{Lift}_{\text{actual}})^2 \\ &= (w^T g(x) - \text{Lift}_{\text{actual}})^2 \\ &= (\sigma(w_1^T \{ \sigma(w_2^T(x)) \}) - \text{Lift}_{\text{actual}})^2 \end{aligned}$$

Now we can see that the Deep Learning model is (in essence) a series of chained basis transformations!

Model

$g(x)$



Why use Deep Learning over other (non-Deep) approaches or not?

Advantages

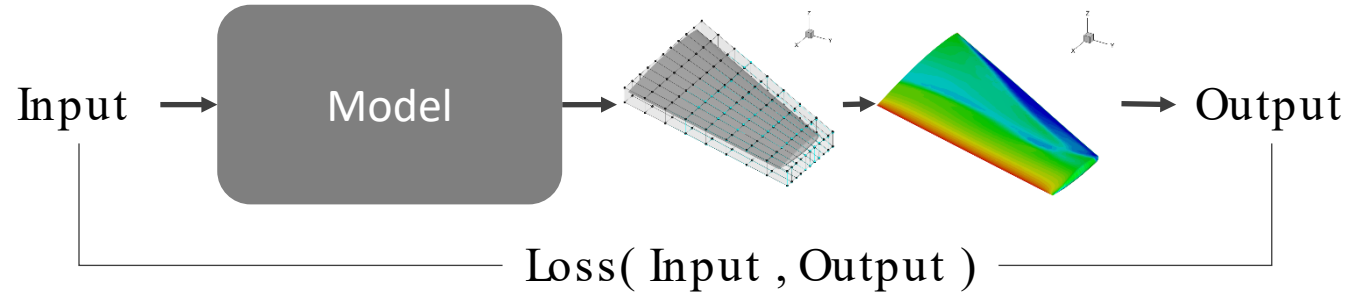
1. Fairly extensible with modern libraries
2. Plays nicely with other differentiable approaches
3. Good hardware acceleration
4. Active research community + industrial investment

Disadvantages


1. Certain modeling assumptions difficult to do
2. Certain architectures have difficulty converging or possess pathologies
3. Theory less developed than some other models

Opportunities and Directions

Merging of Engineering and Deep-Learning models



New invariances & constraints on Deep Learning models

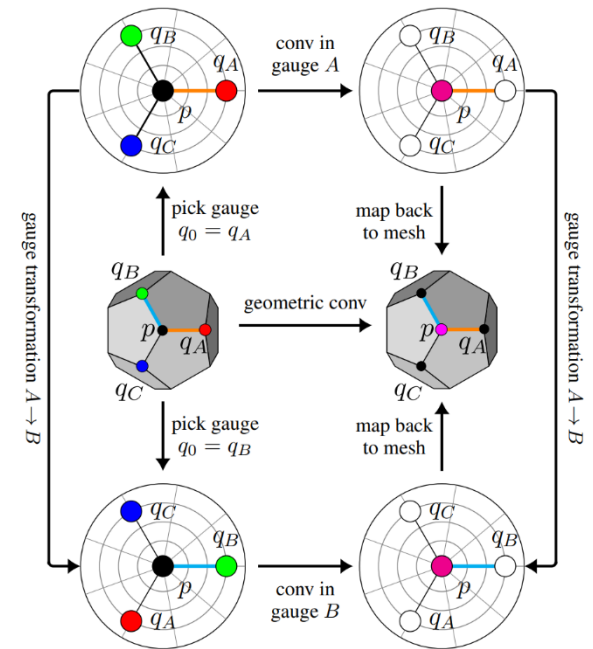


Cohen and Welling, 2016
arxiv.org/abs/1602.07576

Generalizing Convolution

de Haan et al., 2020
arxiv.org/abs/2003.05425

Combining Probabilistic and Deep Learning Models



Thank you

Dr. Mark Fuge

Univ. of Maryland, College Park

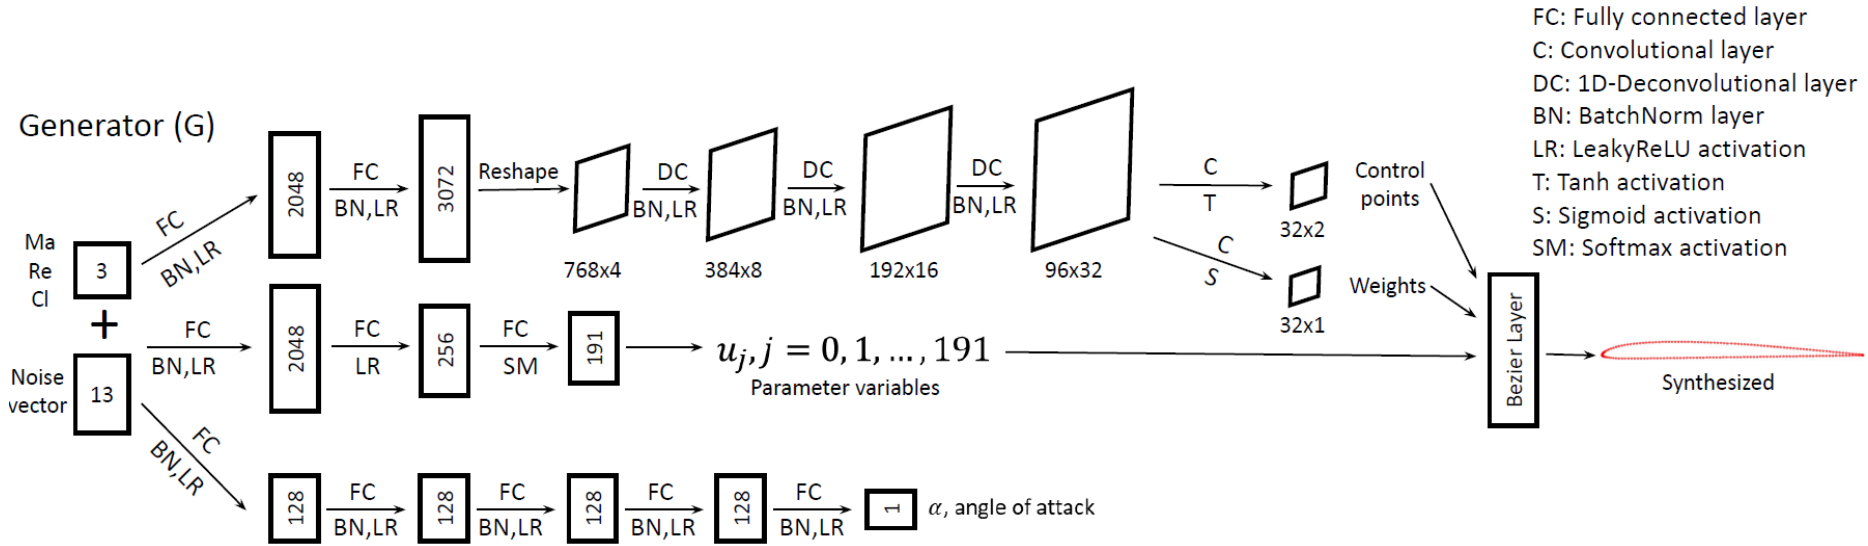
(301) 405-2558

fuge@umd.edu

ideal.umd.edu

Backup Slides

Now we can see that the Deep Learning model is (in essence) a series of chained basis transformations!




Minimize Sinkhorn divergence:

$$\min_G S_\lambda(p_r(\mathbf{x}, \mathbf{y}) || p_g(\mathbf{x}|\mathbf{y})p_r(\mathbf{y}))$$

Loss

<p>Conditional Formulation:</p> $p(\mathbf{x}, \mathbf{y}) = \int_{\hat{\mathbf{y}}, \mathbf{z}} p(\mathbf{x}, \mathbf{y} \hat{\mathbf{y}}, \mathbf{z}) p_r(\hat{\mathbf{y}}) p(\mathbf{z}) d\hat{\mathbf{y}} d\mathbf{z}$ $p(\mathbf{x}, \mathbf{y} \hat{\mathbf{y}}, \mathbf{z}) \propto \exp - \frac{c([\mathbf{x}, \mathbf{y}], [G(\mathbf{z}, \hat{\mathbf{y}}), \hat{\mathbf{y}}])}{\lambda}$	<p>Surrogate Log-Likelihood (SLL):</p> $\log p(\mathbf{x}, \mathbf{y}) \geq -\frac{1}{\lambda} \mathbb{E}_{\mathbb{P}_{Z X,Y}^*} [c([\mathbf{x}, \mathbf{y}], [G^*(\mathbf{z}, \mathbf{y}), \mathbf{y}])] + \mathbb{E}_{\mathbb{P}_Z} [\log p(\mathbf{z})] + H(\mathbb{P}_{Z X,Y}^*) + \log p_r(\mathbf{y}) + \text{const}$ <p>in which</p> $\mathbb{P}_{Z X,Y}^* = \mathbb{P}_Z(\mathbf{z}) \exp \frac{v^*([\mathbf{x}, \mathbf{y}], [G^*(\mathbf{z}, \mathbf{y}), \mathbf{y}])}{\lambda}$	<p>Cost Function:</p> $c([\mathbf{x}, \mathbf{y}, \mathbf{b}], [\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{b}}]) = \mathbf{x} - \hat{\mathbf{x}} + \mathbf{y} - \hat{\mathbf{y}} + \mathbf{b} - \hat{\mathbf{b}} $
--	--	--

What are Generative Models doing?



$$f: \mathcal{Z} \rightarrow \mathcal{X} \quad \mathbb{P}(\mathbf{x}|\mathbf{z})$$

$$f^{-1}: \mathcal{X} \rightarrow \mathcal{Z} \quad \mathbb{P}(\mathbf{z}|\mathbf{x})$$

$$\log \mathbb{P}(\mathbf{x}) = \log \mathbb{P}(\mathbf{z}) + \log |\det \nabla_{\mathbf{x}} f^{-1}(\mathbf{x})|$$