

# Energy-Based Models

## Part 2

# Loss Functions and Architectures

Yann LeCun

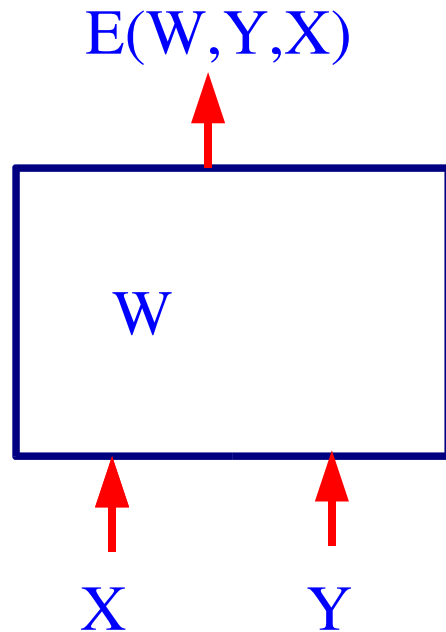
The Courant Institute of Mathematical Sciences

New York University

<http://yann.lecun.com>

<http://www.cs.nyu.edu/~yann>

# Architecture + Inference Algo + Loss Function = Model

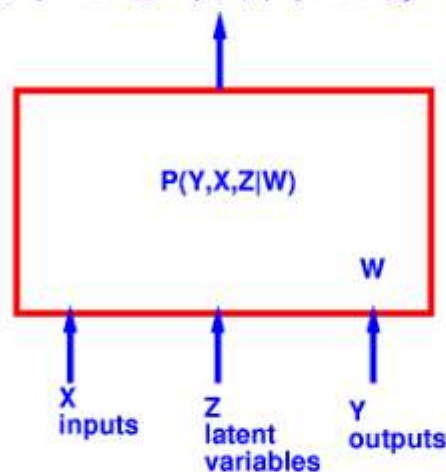


1. **Design an architecture:** a particular form for  $E(W, Y, X)$ .
2. **Pick an inference algorithm for  $Y$ :** MAP or conditional distribution, belief prop, min cut, variational methods, gradient descent, MCMC, HMC.....
3. **Pick a loss function:** in such a way that minimizing it with respect to  $W$  over a training set will make the inference algorithm find the correct  $Y$  for a given  $X$ .
4. **Pick an optimization method.**

**PROBLEM:** What loss functions will make the machine approach the desired behavior?

# Training Probabilistic Models

$$P(Y|X,W) = \text{SUM}_z P(Y,X,z|W) / \text{SUM}_{yz} P(y,X,z|W)$$



Training set:  $\mathcal{S} = \{(X^1, Y^1), \dots, (X^p, Y^p)\}$ .

**Training Criterion:** Max Likelihood

$$\prod_{i=1}^p P(Y^i|X^i, W) = \prod_{i=1}^p \frac{\int_z P(W, Y^i, z, X^i)}{\int_{yz} P(W, y, z, X^i)}$$

**Loss Function:** Negative Log Likelihood:  $\mathcal{L}(W, \mathcal{S}) = -\log \prod_{i=1}^p P(Y^i|X^i, W)$

$$\mathcal{L}(W, \mathcal{S}) = \sum_{i=1}^p -\log (P(W, Y^i, X^i)) + \log \left( \int_y P(W, y, X^i) \right)$$


$$\mathcal{L}(W, \mathcal{S}) = \sum_{i=1}^p -\log \left( \int_z P(W, Y^i, z, X^i) \right) + \log \left( \int_{yz} P(W, y, z, X^i) \right)$$

# What's so bad about probabilistic models?

- Why bother with a normalization since we don't use it for decision making?
- Why insist that  $P(Y|X)$  have a specific shape, when we only care about the position of its minimum?
- When  $Y$  is high-dimensional (or simply combinatorial), normalizing becomes intractable (e.g. Language modeling, image restoration, large DoF robot control...).
- A tiny number of models are pre-normalized (Gaussian, exponential family)
- A very small number are easily normalizable
- A large number have intractable normalization
- A huuuge number can't be normalized at all (examples will be shown).
- Normalization forces us to take into account areas of the space that we don't actually care about because our inference algorithm never takes us there.
- **If we only care about making the right decisions, maximizing the likelihood solves a much more complex problem than we have to.**

# EBM Energy Surfaces

---

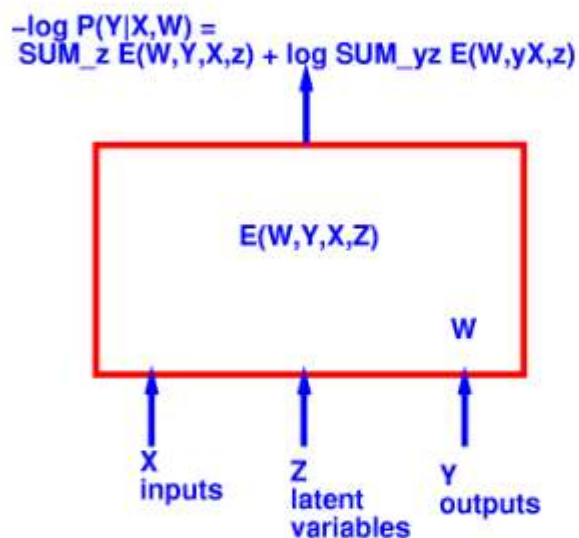


Examples: An EBM that computes  $Y = X^2$ .

**On the left:**  $E(Y, X)$  is quadratic in  $Y$ . It corresponds to a Gaussian model of  $P(Y|X)$ .

**On the right:**  $E(Y, X)$  is saturated. Although it gives the same answers as the EBM on the left, it has no probabilistic equivalent because the partition function  $\int_y \exp(-E(Y, X))$  does not converge.

# Probabilistic Models from Energy-Based Models



- Any joint probability model can be approached as close as we want by an equivalent EBM. If  $P(Y, X, Z|W)$  is non-zero everywhere:  

$$E(W, Y, X, Z) = C - \frac{1}{\beta} \log P(Y, X, Z|W)$$
 where  $C$  is an arbitrary constant and  $\beta$  a strictly positive constant.
- not all EBMs can be turned into a probabilistic model. Only those for which  $\int_{yz} \exp(-\beta E(W, y, X, z))$  converges:

$$P(Y|X) = \frac{\int_z \exp(-\beta E(W, Y, X, z))}{\int_{yz} \exp(-\beta E(W, y, X, z))}$$

Any single probabilistic model will have many equivalent EBMs when it comes to comparison-based inference or decision. *Because many energy surfaces have minima at the same places.*

**We have a lot more flexibility with EBMs than with Prob. Models**

# What's good/bad about EBM?

---

## What's bad about EBMs:

- There is no compositionality...
- ... but we don't care because we are going to train our whole system *end-to-end*.  
With *end-to-end learning*, we do not need compositionality.

## What's good about EBMs:

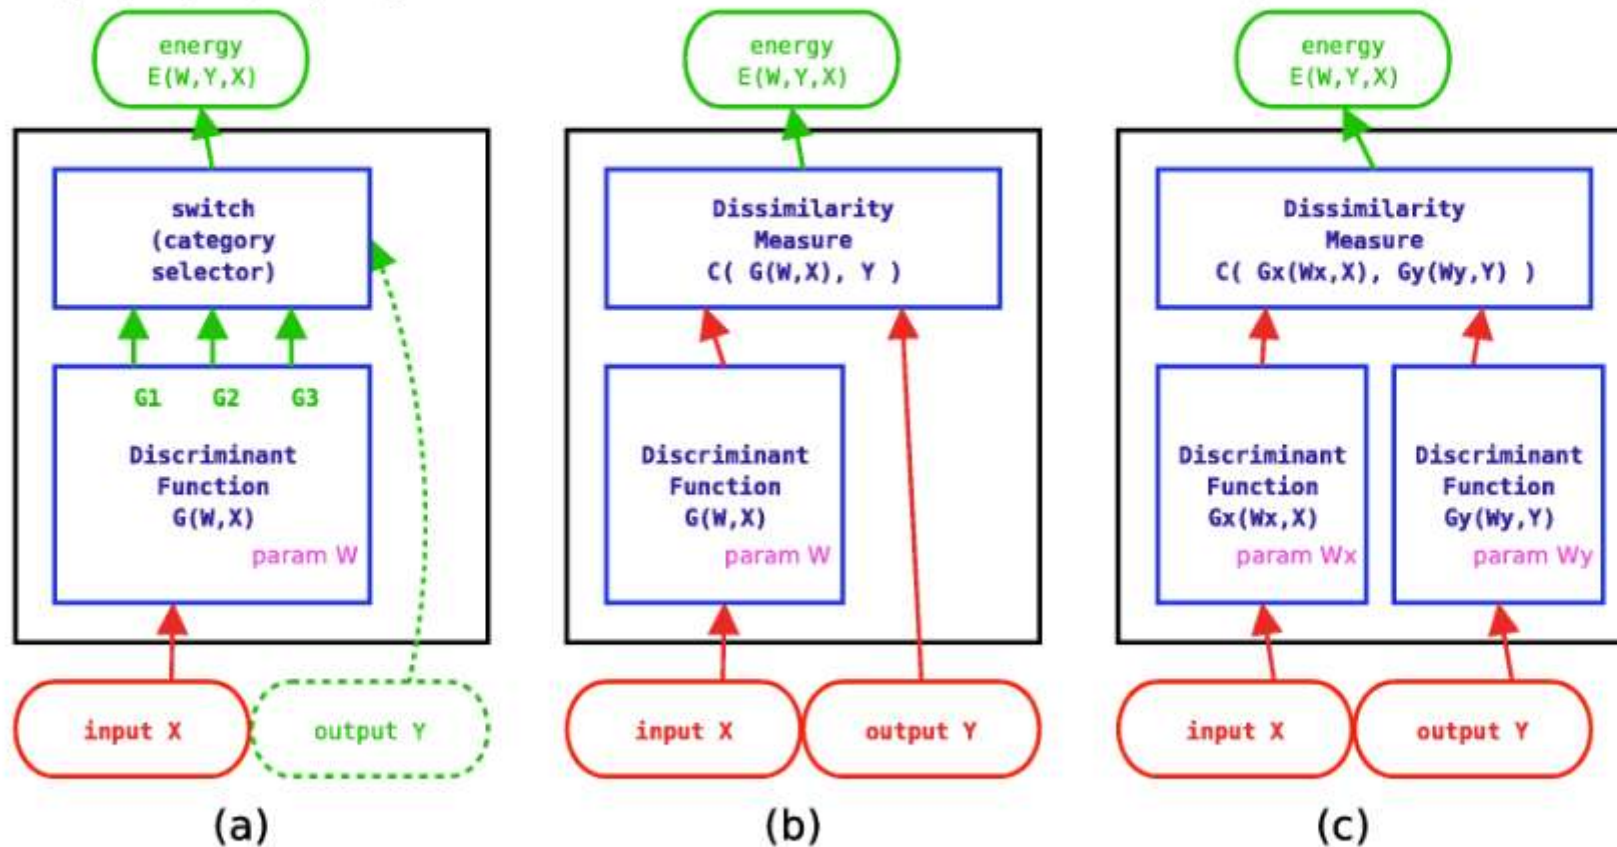
- We have complete freedom for the form and parameterization of the energy function (including things that can't be normalized).
- because we do not need to normalize, we can use a much larger repertoire of model architectures.
- No need for computing (intractable) partition functions
- No need to justify the choice of your favorite approximation of the partition function.

Pretty much every model we know is some form of EBM.

**QUESTION:** what loss functions can we use for training?

# Examples of EBM

Almost every type of model we know is some form of EBM. It all depends on how  $E(W, Y, X, Z)$  is parameterized. NOTE: there is no linearity assumption.





# Training EBMs

---

- Training will consist in finding a  $W$  that minimizes a **loss function**  $\mathcal{L}(W, \mathcal{S})$ , over the training set  $\mathcal{S}$ .
- We must devise loss functions that **“carve”** the energy landscape so that the energy is **small around training samples** and **high everywhere else..**
- We seek loss functions that **do not require** evaluating intractable integrals, but which, nevertheless, drive the machine to approach the desired behavior.
- Basic idea: **“dig holes”** at  $(X, Y)$  locations near training samples, while **“building hills”** at un-desired locations, particularly the ones that are erroneously picked by the inference algorithm.
- Whereas probabilistic models trained with max likelihood shape the entire energy surface, our EBM loss function will merely dig holes at the right places and build hills only where needed to avoid erroneous inferences.

# Loss Functions for EBMs

---

- training set  $\mathcal{S} = \{(X^i, Y^i), i = 1..p\}$

- Loss:

$$\mathcal{L}(W, \mathcal{S}) = R \left( \frac{1}{p} \sum_{i=1}^p L(W, Y^i, X^i) \right)$$

- $L(W, Y^i, X^i)$  is the per-sample loss function for sample  $(X^i, Y^i)$ .  $L$  is assumed to have a lower bound.
- $R$  is a monotonically increasing function. In the following we assume  $R$ =identity
- the loss is invariant under permutations of the samples, and under multiple repetitions of the same training set.
- What form can  $L(W, Y, X)$  take?

## Conditions on the Loss

---

- Condition for correct output on sample  $(X^i, Y^i)$ : there is a margin  $m > 0$ , such that:

$$\check{E}(W, Y^i, X^i) < \check{E}(W, Y, X^i) - m, \quad \forall Y \in \{Y\}, Y \neq Y^i$$

- **Assumption:**  $L$  depends on  $X^i$  only through the set of energies  $\{\check{E}(W, Y, X^i), Y \in \{Y\}\}$ .

- For example, if  $\{Y\} = \{0, 1, \dots, k - 1\}$

$$L(W, Y^i, X^i) = L(Y^i, \check{E}(W, 0, X^i), \dots, \check{E}(W, k - 1, X^i))$$

- We want to design  $L$  so that making an update of  $W$  to decrease  $L(W, Y^i, X^i)$  will automatically decrease the difference  $\check{E}(W, Y^i, X^i) - \check{E}(W, Y, X^i)$  for values of  $Y$  such that  $\check{E}(W, Y^i, X^i) < \check{E}(W, Y, X^i) - m$ .

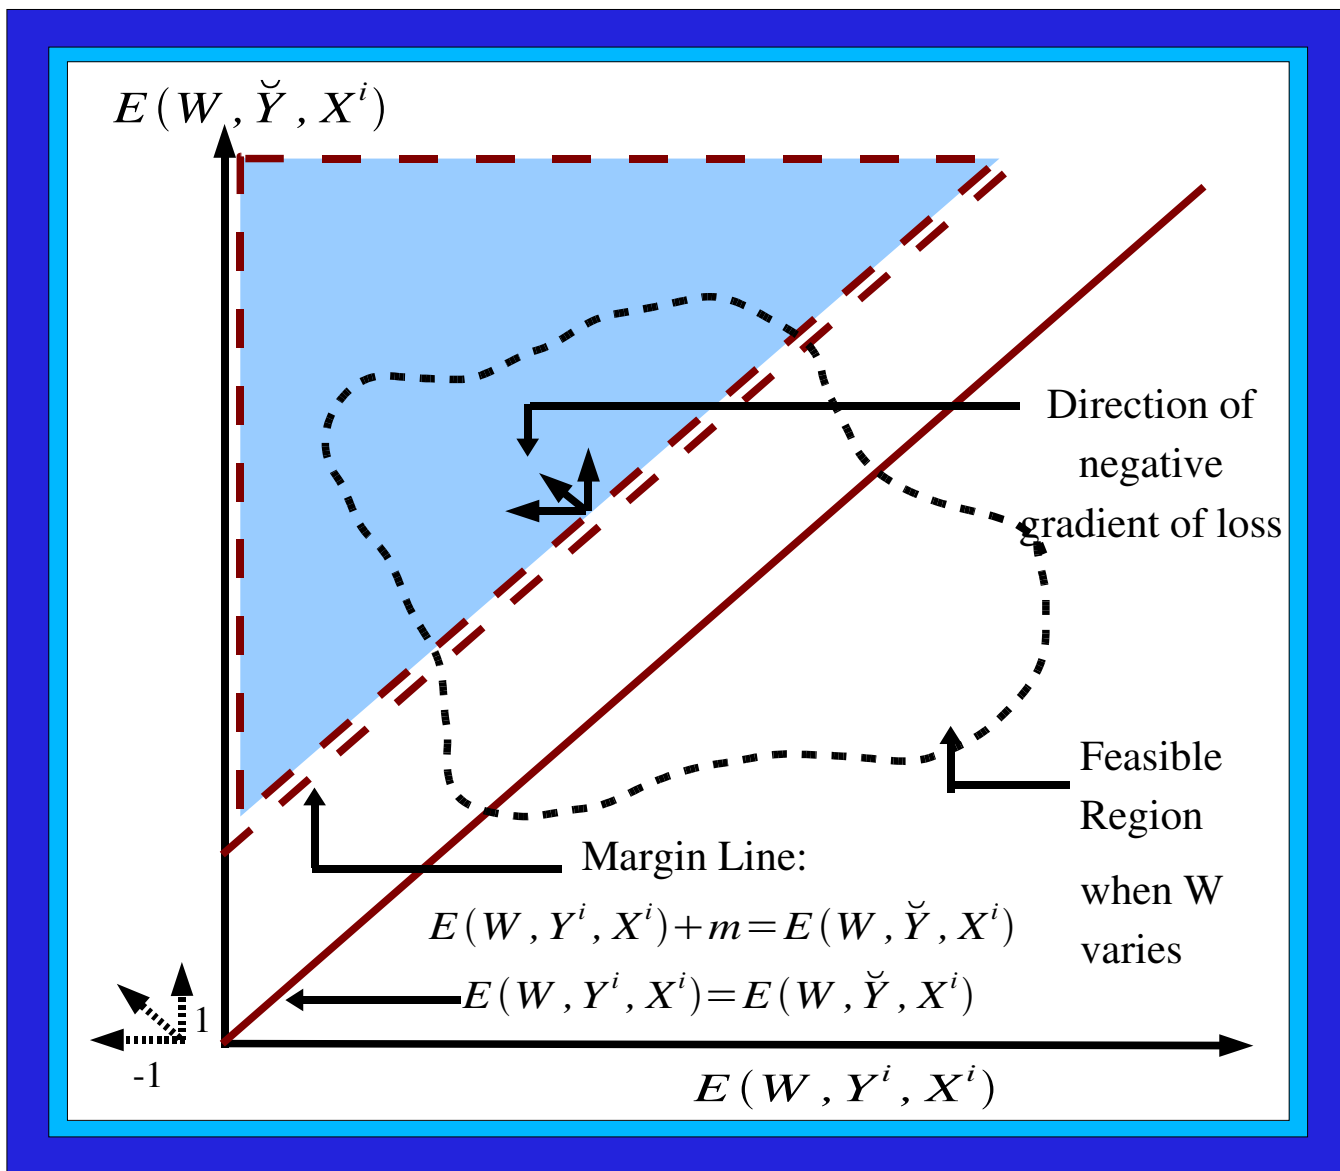
## Conditions on the Loss

- Let's define  $\bar{Y}$  as the most offending incorrect output:

$$\bar{Y} = \operatorname{argmin}_{y \in \{Y\}, y \neq Y^i} \check{E}(W, y, X^i)$$

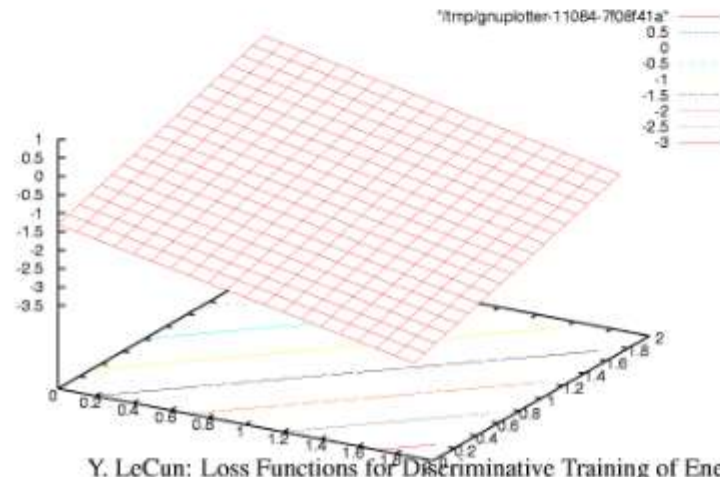



- **Cond. for correct output:**  $\check{E}(W, Y^i, X^i) < \check{E}(W, \bar{Y}, X^i) - m$
- Let's assume that  $L(W, Y^i, X^i)$  is convex in the 2 coordinates  $\check{E}(W, Y^i, X^i)$  and  $\check{E}(W, \bar{Y}, X^i)$ .
- **Sufficient condition 1:** All the minima of  $L(W, Y^i, X^i)$  must be in the half-plane  $\check{E}(W, Y^i, X^i) < \check{E}(W, \bar{Y}, X^i) - m$ .
- **Sufficient condition 2:** The gradient of  $L(W, Y^i, X^i)$  on the margin line  $\check{E}(W, Y^i, X^i) = \check{E}(W, \bar{Y}, X^i) - m$ , must have a positive dot product with the direction  $[-1, 1]$ .
- **Sufficient condition 3:** On the margin line  $\check{E}(W, \bar{Y}, X^i) = \check{E}(W, Y^i, X^i) + m$ , the following must hold:  $\left[ \frac{\partial \check{E}(W, Y^i, X^i)}{\partial W} - \frac{\partial \check{E}(W, \bar{Y}, X^i)}{\partial W} \right] \cdot \frac{\partial L(W, Y^i, X^i)}{\partial W} > 0$

# Conditions on the Loss



# Condition on the Loss

Loss  $L(W, Y^i, X^i)$  as a function of  $\tilde{E}(W, \bar{Y}, X^i)$  and  $\check{E}(W, Y^i, X^i)$



# Examples of Loss Functions

---

- **Energy Loss:**  $L_{\text{energy}}(W, Y^i, X^i) = \check{E}(W, Y^i, X^i)$ .

Only works if the architecture is such that decreasing  $\check{E}(W, Y^i, X^i)$  will automatically increase  $\check{E}(W, Y, X^i)$  for  $y \neq Y^i$ .

- **Generalized Perceptron Loss** [LeCun 1998][Collins 2002]:

$$L_{\text{ptron}}(W, Y^i, X^i) = \check{E}(W, Y^i, X^i) - \min_{Y \in \{Y\}} \check{E}(W, Y, X^i)$$

Does not work because the margin is zero. This reduces to the traditional linear perceptron loss when  $\check{E}(W, Y, X) = -YW.X$ .

- **Generalized Margin Loss:** [LeCun et al. 2005]

$$L_{\text{gmargin}}(W, Y^i, X^i) = Q[\check{E}(W, Y^i, X^i), \check{E}(W, \bar{Y}, X^i)]$$

Where  $Q$  is an increasing function of  $\check{E}(W, Y^i, X^i)$  and a decreasing function of  $\check{E}(W, \bar{Y}, X^i)$ .


- **Negative Log Likelihood Loss:** [Bengio 1992][LeCun 1998][Lafferty 2001]

$$L_{\text{nll}}(W, Y^i, X^i) = \check{E}(W, Y^i, X^i) - F_{\beta}(W, X^i)$$

$$\text{with: } F_{\beta}(W, X^i) = -\frac{1}{\beta} \log \left( \int_{Y \in \{Y\}} \exp[-\beta \check{E}(W, Y, X^i)] \right)$$


# Special Cases of the Generalized Margin Loss

$$L_{\text{gmargin}}(W, Y^i, X^i) = Q[\check{E}(W, Y^i, X^i), \check{E}(W, \bar{Y}, X^i)]$$




**Hinge Loss** [Taskar, Guestrin, Koller 2003],[Altun, Johnson, Hofmann, 2003]:

$$L_{\text{hinge}}(W, Y^i, X^i) = \max(0, m + \check{E}(W, Y^i, X^i) - \check{E}(W, \bar{Y}, X^i))$$



**Square-Square Loss:**

$$L_{\text{sqsq}}(W, Y^i, X^i) = \check{E}(W, Y^i, X^i)^2 + (\min(0, m - \check{E}(W, \bar{Y}, X^i)))^2$$




**Square-Exp Loss:** [Osadchy, Miller, LeCun, NIPS 2004]

$$L_{\text{sqexp}}(W, Y^i, X^i) = \check{E}(W, Y^i, X^i)^2 + K \exp(-\beta \check{E}(W, \bar{Y}, X^i))$$



# EBM Demos



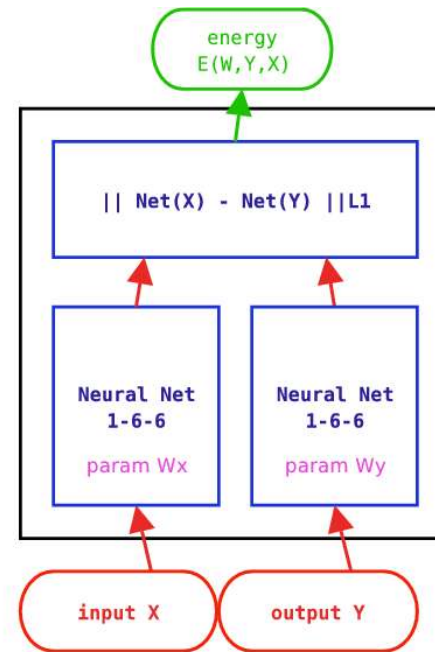
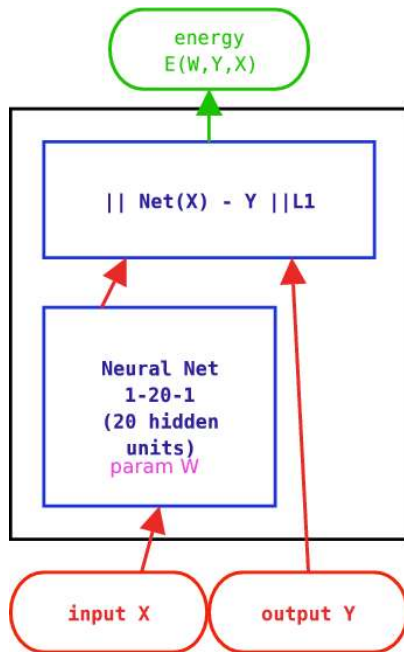
- Demo 1:  $Y = X^2$ , Architecture A, Square Energy Loss. It works because  $E(Y, X)$  is a fixed quadratic function of  $Y$ .
- Demo 2:  $Y = X^2$ , Architecture B, Square Energy. It collapses.
- Demo 3:  $Y = X^2$ , Architecture B, Square-Square Margin Loss
- Demo 4:  $Y = X^2$ , Architecture B, Negative Log Likelihood Loss. Few iterations, but each iteration is expensive

Initially, the forbidden sphere around  $Y^i$  is 0.2, then 0.1.



Demo 5: eye pattern, Architecture B, Negative Log Likelihood Loss.

# EBM Demos: energy loss

Loss: “Energy Loss”:  $L(W, Y, X) = E(W, Y, X)$

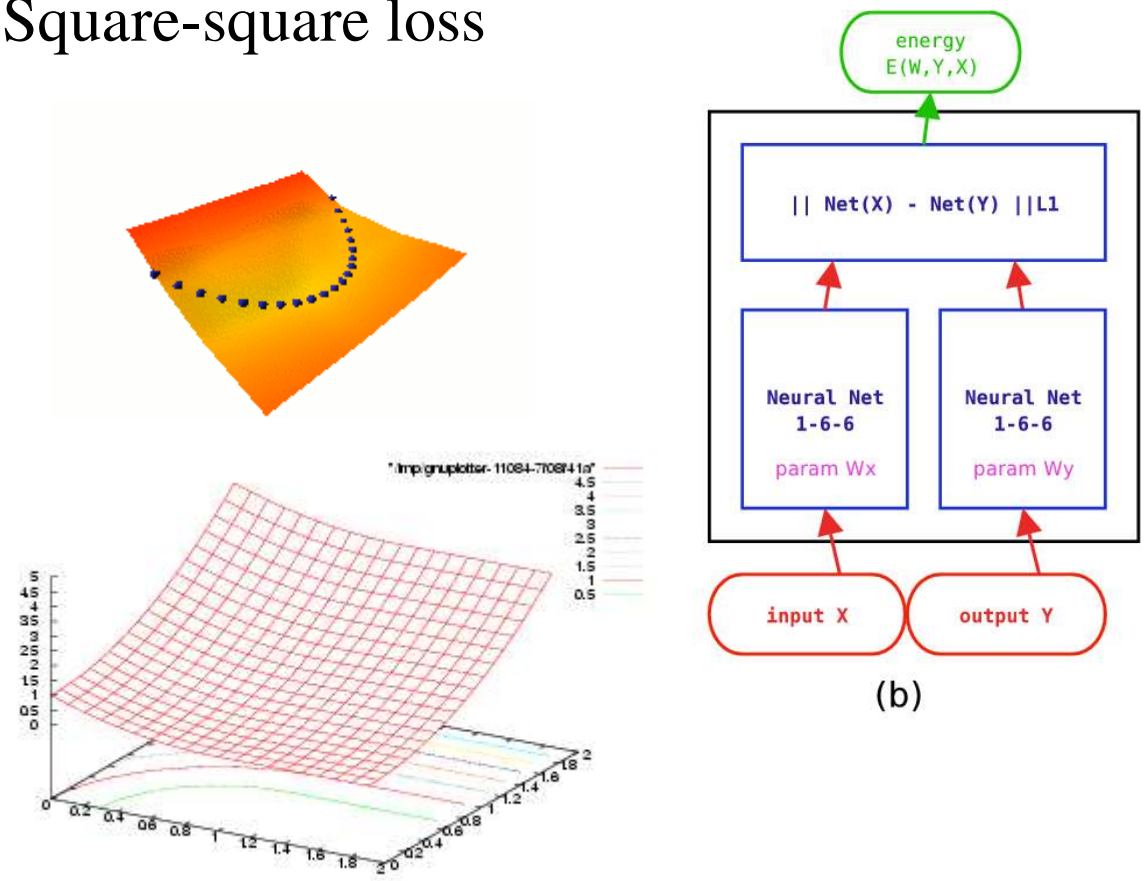


**COLLAPSE!!!**



# EBM Demos: good losses

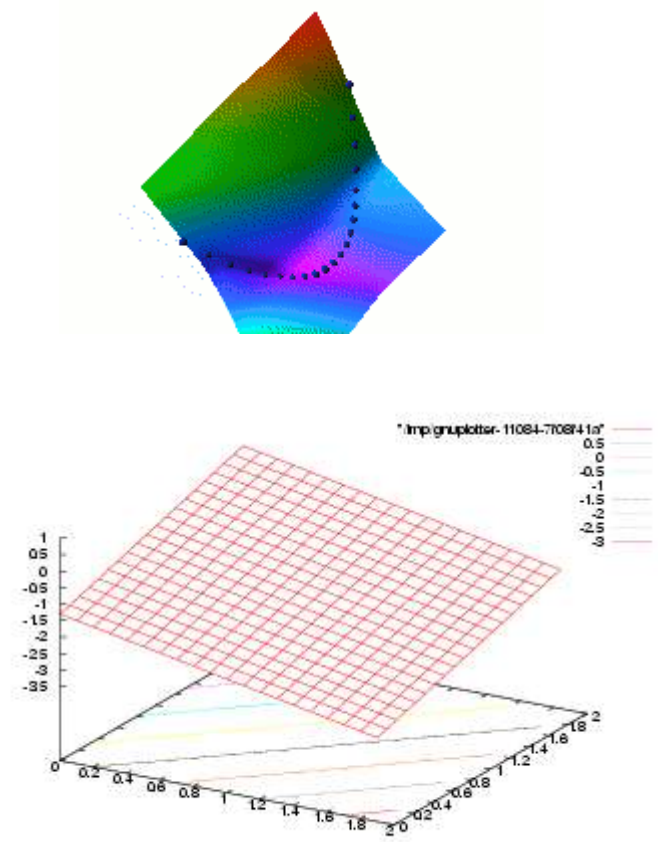
## Square-square loss



$$L(W, X, Y) = E(W, Y, X)^2 - \max(0, m - E(W, \bar{Y}, X)^2)$$

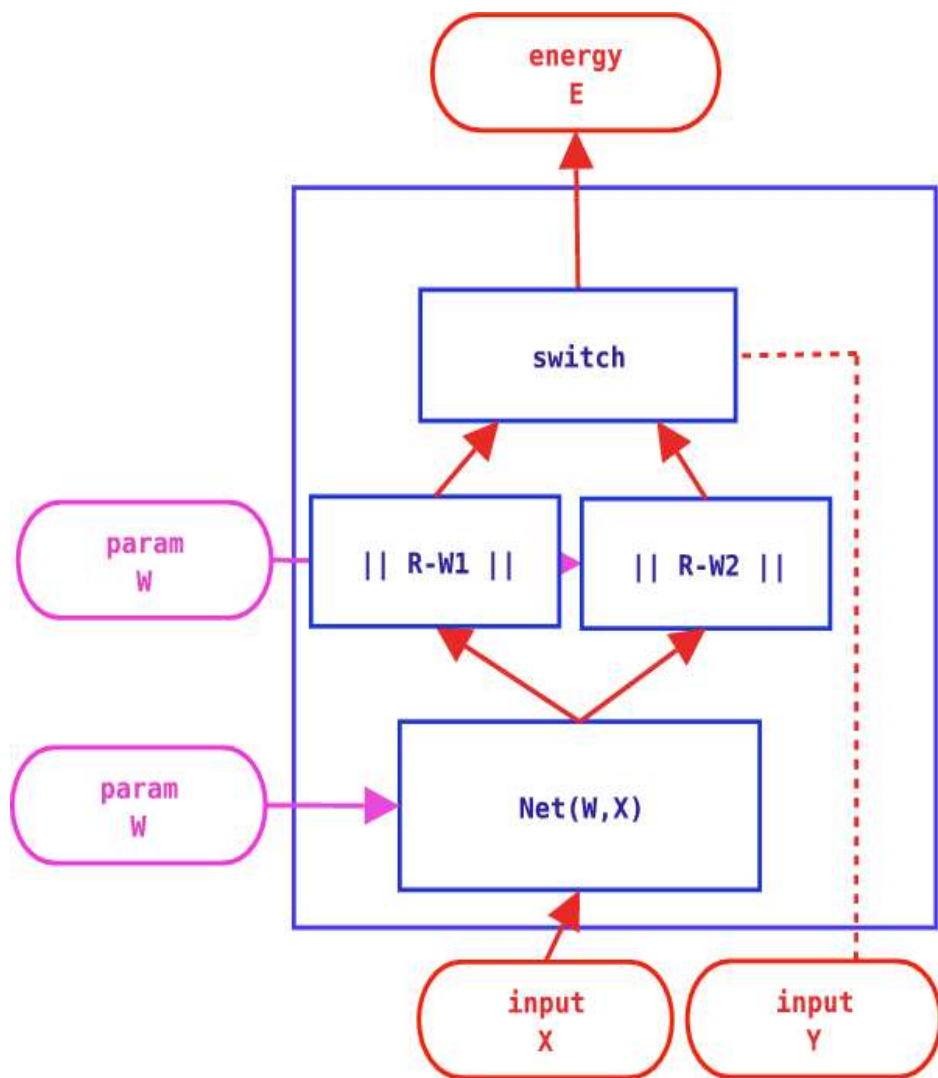
$$\bar{Y} = \operatorname{argmin}_{y \in \{Y\}, y \neq Y} E(W, y, X)$$

## Neg-Log-Likelihood loss



$$L(W, X, Y) = E(W, Y, X) + \frac{1}{\beta} \log \left[ \sum_{y \in \{Y\}} \exp(-\beta E(W, y, X)) \right]$$

## Other Architectures that may collapse



### Linear module followed by radial basis functions and a switch:

- ▶ Will collapse with the energy loss and the perceptron loss.
- ▶ Will not collapse with the square-square, neg-log-likelihood, margin loss, etc....


# EBM

- Unlike traditional classifiers, EBMs can represent **multiple alternative outputs**
- The normalization in probabilistic models is often an unnecessary aggravation, particularly if the ultimate goal of the system is to make decisions.
- EBMs with appropriate loss function avoid the necessity to compute the partition function and its derivatives (which may be intractable)
- EBMs give us complete freedom in the choice of the architecture that models the joint “incompatibility” (energy) between the variables.
- We can use architectures that are not normally allowed in the probabilistic framework (like neural nets).
- **The inference algorithm that finds the most offending (lowest energy) incorrect answer does not need to be exact:** our model may give **low energy** to far-away regions of the landscape. But if our inference algorithm **never finds those regions, they do not affect us.** But **they do affect normalized probabilistic models**

# Non-Linear Models

- We can always add “features”, “kernels”, or “basis function” in front of a linear model to make it non-linear.
  - ▶ This is how non-linear SVMs are built.
- **Question:** so, why would we need anything else?
- **Answer:** the complexity of the real world is very difficult to capture in a kernel.
- **How do we solve:**
  - ▶ The invariance problem in image recognition.
  - ▶ The structured output problem in sequence labeling (parts of speech tagging, speech recognition, biological sequence analysis....).

# Fixed Preprocessing (features, kernels, basis functions)



- Map the inputs into a (higher dimensional) “feature” space
  - ▶ With more dimensions, the task is more likely to be linearly separable.
- **Problem:** how should we pick the features so that the task becomes linearly separable in the feature space?
  - ▶ **Classical approach 1:** we use our prior knowledge about the problem to hand-craft an appropriate feature set.
  - ▶ **Classical Approach 2:** we use a “standard” set of basis functions (RBFs....)


# Fixed Preprocessing: Kernels and SVMs

## Simplest approach: The Kernel Method (thanks to Wahba's Representer Theorem)

- ▶ Make each basis function a “bump” function (a template matcher).
- ▶ Place one bump around each training sample.
- ▶ Compute a linear combination of the bumps.
- ▶ In the “bump space”, we get one separate dimension for each training sample, so if the bumps are narrow enough, we can learn any mapping on the training set.
- ▶ To generalize on unseen samples, we adjust the bump widths and we regularize the weights.
- ▶ We get a **Support Vector Machine**.

**Problem:** an SVM is a glorified template matcher which is only as good as its kernel.

$$G(X, \alpha) = \sum_j \alpha_j K(X^j, X)$$





# Trainable Front-End, Structured Architectures

## • The Solutions:

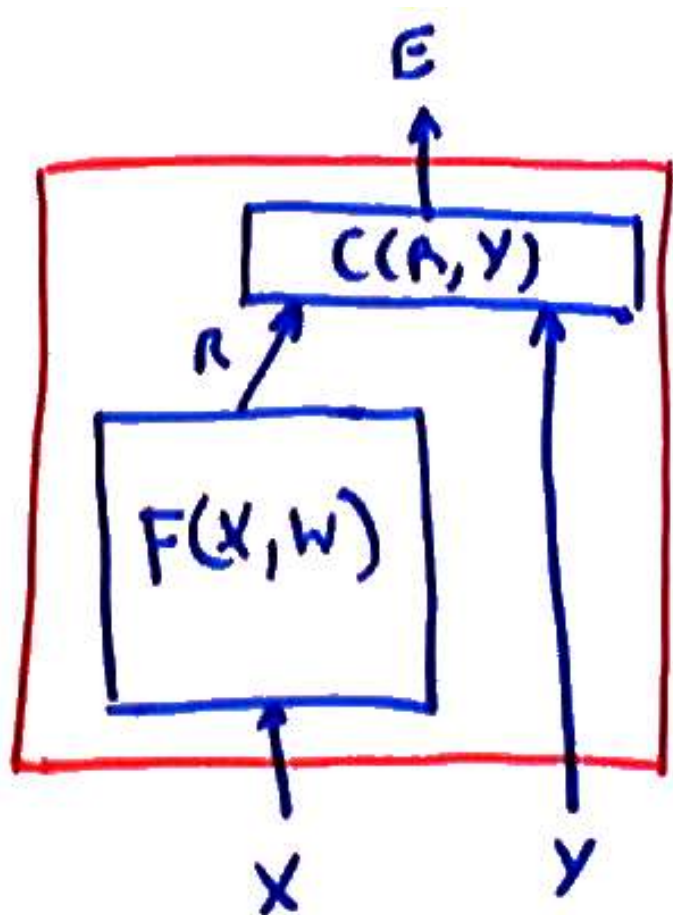
- ▶ **Invariance:** Do not use a fixed front-end, **make it trainable**, so it can learn to extract invariant representations
- ▶ **Structure:** Do not use simple linearly-parameterized classifiers, use architectures whose inference process involves multiple non-linear decisions, as well as search and “reasoning”.

## • We need total flexibility in the design of the architecture of the machine:

- ▶ So that we can tailor the architecture to the task
- ▶ So that we can build our prior knowledge about the task into the architecture

## • Multi-Module Architectures.

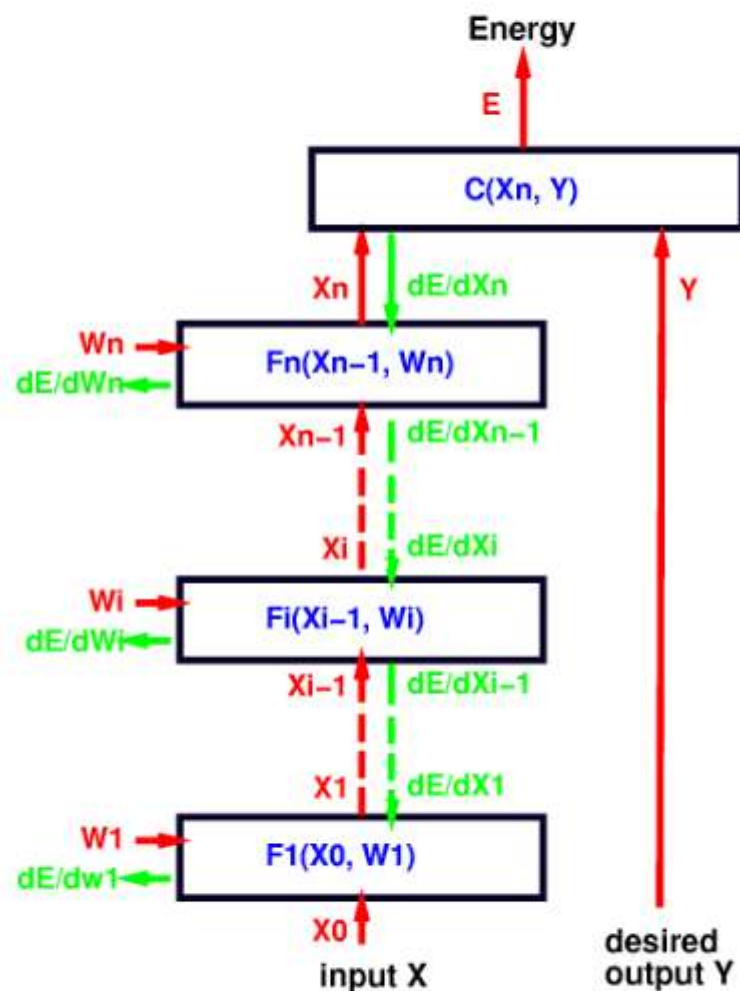
# Multi-Module Architectures



## • For Supervised Learning

- ▶ We allow the function  $F(W, X)$  to be non-linearly parameterized in  $W$ .
- ▶ This allows us to play with a large repertoire of functions with rich class boundaries.
- ▶ We assume that  $F(W, X)$  is differentiable almost everywhere with respect to  $W$ .

# Multi-Module Systems: Cascade

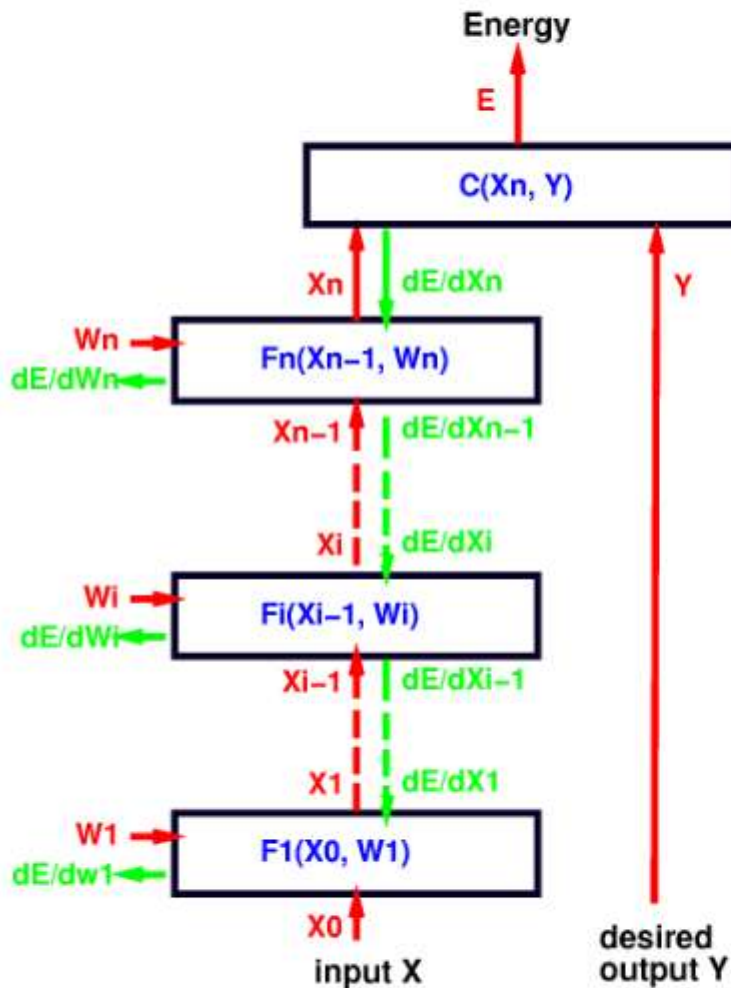


- Complex learning machines can be built by assembling Modules into networks.
- a simple example: layered, feed-forward architecture (cascade).
- computing the output from the input:  
**forward propagation**
- let  $X = X_0$ ,

$$X_i = F_i(X_{i-1}, W_i) \quad \forall i \in [1, n]$$

$$E(Y, X, W) = C(X_n, Y)$$

# Object-Oriented Implementation



- Each module is an object (instance of a class).
- Each class has an “fprop” (forward propagation) method that takes the input and output states as arguments and computes the output state from the input state.
- Lush:  
(==> module fprop input output)
- C++:  
`module.fprop(input, output);`

# Gradient of the Loss, gradient of the Energy

---

- We assumed early on that the loss depends on  $W$  only through the terms  $E(W, Y, X^i)$ :

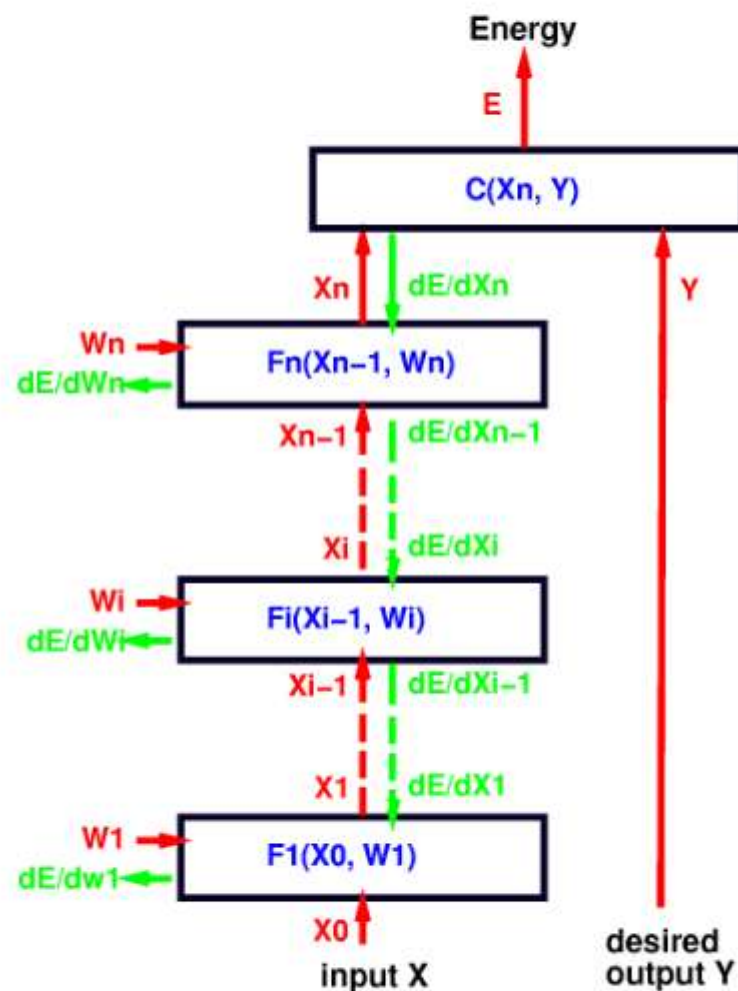
$$L(W, Y^i, X^i) = L(Y^i, E(W, 0, X^i), E(W, 1, X^i), \dots, E(W, k-1, X^i))$$

- therefore:

$$\frac{\partial L(W, Y^i, X^i)}{\partial W} = \sum_Y \left[ \frac{\partial L(W, Y^i, X^i)}{\partial E(W, Y, X^i)} \frac{\partial E(W, Y, X^i)}{\partial W} \right]$$

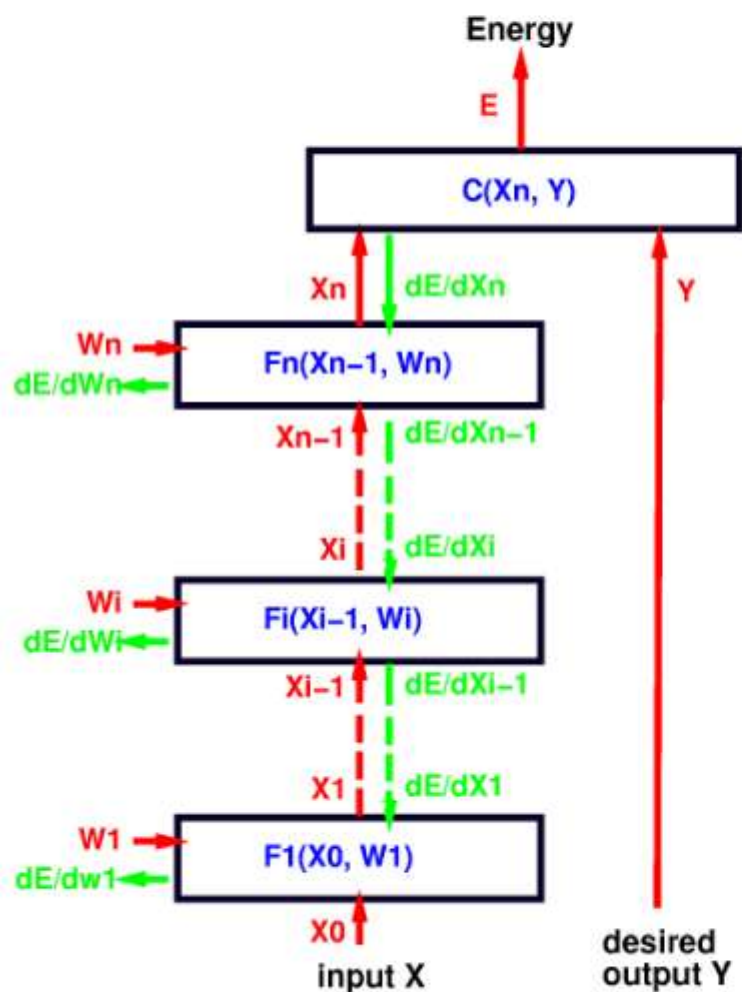
- We only need to compute the terms  $\frac{\partial E(W, Y, X^i)}{\partial W}$
- Question: How do we compute those terms efficiently?

# Computing the Gradients in Multi-Layer Systems



- To train a multi-module system, we must compute the gradient of  $E$  with respect to all the parameters in the system (all the  $W_i$ ).
- Let's consider module  $i$  whose forward method computes  $X_i = F_i(X_{i-1}, W_i)$ .
- Let's assume that we already know  $\frac{\partial E}{\partial X_i}$ , in other words, for each component of vector  $X_i$  we know how much  $E$  would wiggle if we wiggled that component of  $X_i$ .

# Computing the Gradients in Multi-Layer Systems



- We can apply chain rule to compute  $\frac{\partial E}{\partial W_i}$  (how much  $E$  would wiggle if we wiggled each component of  $W_i$ ):

$$\frac{\partial E}{\partial W_i} = \frac{\partial E}{\partial X_i} \frac{\partial F_i(X_{i-1}, W_i)}{\partial W_i}$$

$$[1 \times N_w] = [1 \times N_x] \cdot [N_x \times N_w]$$

- $\frac{\partial F_i(X_{i-1}, W_i)}{\partial W_i}$  is the *Jacobian matrix* of  $F_i$  with respect to  $W_i$ .

$$\left[ \frac{\partial F_i(X_{i-1}, W_i)}{\partial W_i} \right]_{kl} = \frac{\partial [F_i(X_{i-1}, W_i)]_k}{\partial [W_i]_l}$$


- Element  $(k, l)$  of the Jacobian indicates how much the  $k$ -th output wiggles when we wiggle the  $l$ -th weight.

# Computing the Gradients in Multi-Layer Systems

Using the same trick, we can compute  $\frac{\partial E}{\partial X_{i-1}}$ . Let's assume again that we already know  $\frac{\partial E}{\partial X_i}$ , in other words, for each component of vector  $X_i$  we know how much  $E$  would wiggle if we wiggled that component of  $X_i$ .

- We can apply chain rule to compute  $\frac{\partial E}{\partial X_{i-1}}$  (how much  $E$  would wiggle if we wiggled each component of  $X_{i-1}$ ):

$$\frac{\partial E}{\partial X_{i-1}} = \frac{\partial E}{\partial X_i} \frac{\partial F_i(X_{i-1}, W_i)}{\partial X_{i-1}}$$



- $\frac{\partial F_i(X_{i-1}, W_i)}{\partial X_{i-1}}$  is the *Jacobian matrix* of  $F_i$  with respect to  $X_{i-1}$ .
- $F_i$  has two Jacobian matrices, because it has two arguments.
- Element  $(k, l)$  of this Jacobian indicates how much the  $k$ -th output wiggles when we wiggle the  $l$ -th input.
- **The equation above is a recurrence equation!**



# Jacobians and Dimensions

- derivatives with respect to a column vector are line vectors (dimensions:  $[1 \times N_{i-1}] = [1 \times N_i] * [N_i \times N_{i-1}]$ )

$$\frac{\partial E}{\partial X_{i-1}} = \frac{\partial E}{\partial X_i} \frac{\partial F_i(X_{i-1}, W_i)}{\partial X_{i-1}}$$

- (dimensions:  $[1 \times N_{wi}] = [1 \times N_i] * [N_i \times N_{wi}]$ ):

$$\frac{\partial E}{\partial W_i} = \frac{\partial E}{\partial X_i} \frac{\partial F_i(X_{i-1}, W_i)}{\partial W}$$

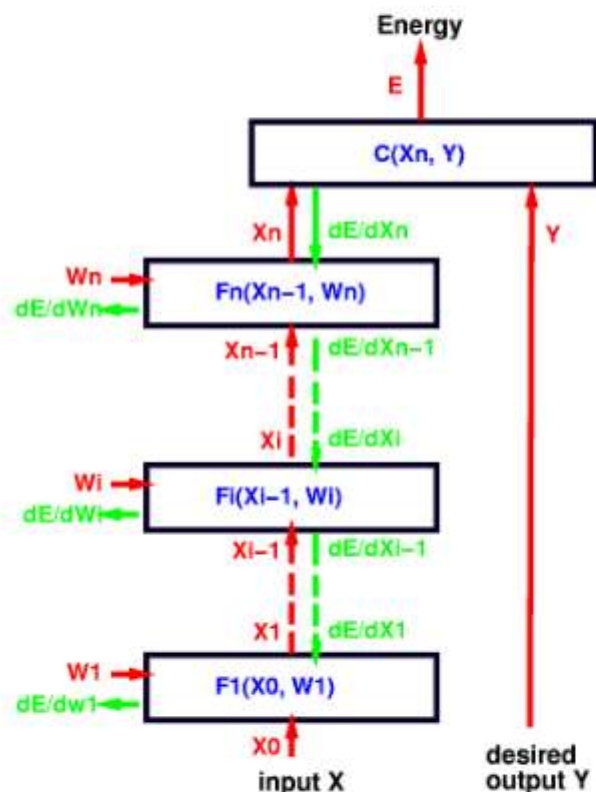
- we may prefer to write those equation with column vectors:

$$\frac{\partial E}{\partial X_{i-1}}' = \frac{\partial F_i(X_{i-1}, W_i)'}{\partial X_{i-1}} \frac{\partial E}{\partial X_i}'$$

$$\frac{\partial E}{\partial W_i}' = \frac{\partial F_i(X_{i-1}, W_i)'}{\partial W} \frac{\partial E}{\partial X_i}'$$

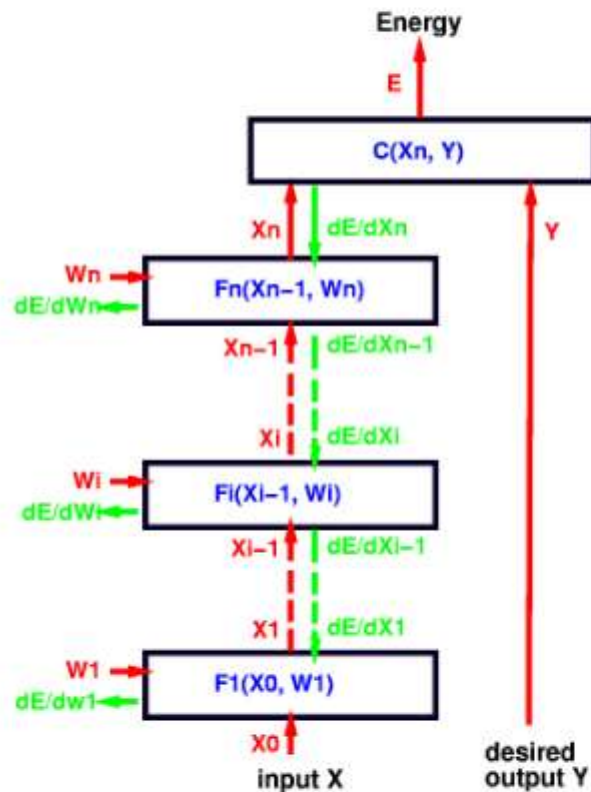
# Back-propagation

To compute all the derivatives, we use a backward sweep called the **back-propagation algorithm** that uses the recurrence equation for  $\frac{\partial E}{\partial X_i}$



- $\frac{\partial E}{\partial X_n} = \frac{\partial C(X_n, Y)}{\partial X_n}$
- $\frac{\partial E}{\partial X_{n-1}} = \frac{\partial E}{\partial X_n} \frac{\partial F_n(X_{n-1}, W_n)}{\partial X_{n-1}}$
- $\frac{\partial E}{\partial W_n} = \frac{\partial E}{\partial X_n} \frac{\partial F_n(X_{n-1}, W_n)}{\partial W_n}$
- $\frac{\partial E}{\partial X_{n-2}} = \frac{\partial E}{\partial X_{n-1}} \frac{\partial F_{n-1}(X_{n-2}, W_{n-1})}{\partial X_{n-2}}$
- $\frac{\partial E}{\partial W_{n-1}} = \frac{\partial E}{\partial X_{n-1}} \frac{\partial F_{n-1}(X_{n-2}, W_{n-1})}{\partial W_{n-1}}$
- ....etc, until we reach the first module.
- we now have all the  $\frac{\partial E}{\partial W_i}$  for  $i \in [1, n]$ .

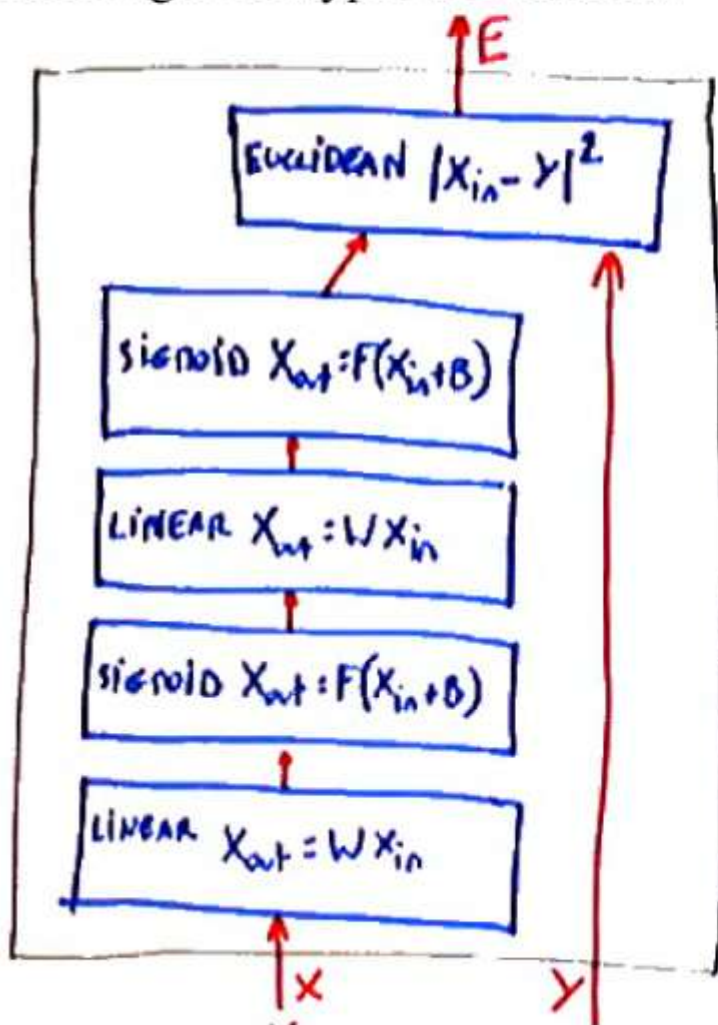
# Object-Oriented Implementation



- Each module is an object (instance of a class).
- Each class has a “bprop” (backward propagation) method that takes the input and output states as arguments and computes the derivative of the energy with respect to the input from the derivative with respect to the output:
  - Lush: (`==>` module bprop input output)
  - C++: `module.bprop(input, output);`
  - the objects input and output contain two slots: one vector for the forward state, and one vector for the backward derivatives.
  - the method `bprop` computes the backward derivative slot of input, by multiplying the backward derivative slot of output by the Jacobian of the module at the forward state of input.

# Modules in a Multi-layer Neural Net

A fully-connected, feed-forward, multi-layer neural nets can be implemented by stacking three types of modules.

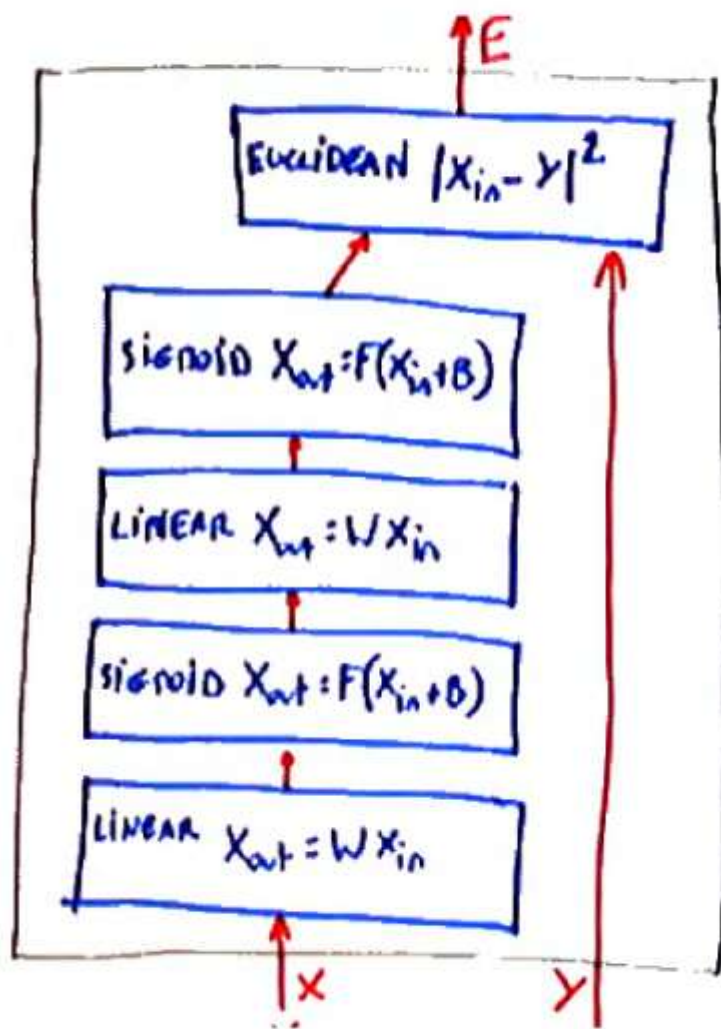


- Linear modules:  $X_{in}$  and  $X_{out}$  are vectors, and  $W$  is a weight matrix.

$$X_{out} = W X_{in}$$

- Sigmoid modules:  
 $(X_{out})_i = \sigma((X_{in})_i + B_i)$  where  $B$  is a vector of trainable “biases”, and  $\sigma$  is a sigmoid function such as tanh or the logistic function.
- a Euclidean Distance module  $E = \frac{1}{2} \|Y - X_{in}\|^2$ . With this energy function, we will use the neural network as a regressor rather than a classifier.

# Loss Function




Here, we will use the simple Energy Loss function  $L_{\text{energy}}$ :

$$L_{\text{energy}}(W, Y^i, X^i) = E(W, Y^i, X^i)$$

## OO Implementation: the `state1` Class


---



the internal state of the network will be kept in a “state” class that contains two scalars, vectors, or matrices: (1) the state proper, (2) the derivative of the energy with respect to that state.

# Linear Module

The input vector is multiplied by the weight matrix.



- fprop:  $X_{out} = W X_{in}$
- bprop to input:  
$$\frac{\partial E}{\partial X_{in}} = \frac{\partial E}{\partial X_{out}} \frac{\partial X_{out}}{\partial X_{in}} = \frac{\partial E}{\partial X_{out}} W$$
- by transposing, we get column vectors:  
$$\frac{\partial E}{\partial X_{in}}' = W' \frac{\partial E}{\partial X_{out}}'$$
- bprop to weights:  
$$\frac{\partial E}{\partial W_{ij}} = \frac{\partial E}{\partial X_{outi}} \frac{\partial X_{outi}}{\partial W_{ij}} = X_{in j} \frac{\partial E}{\partial X_{outi}}$$
- We can write this as an outer-product:  
$$\frac{\partial E}{\partial W}' = \frac{\partial E}{\partial X_{out}}' X_{in}'$$

# Linear Module


---

Lush implementation:

```
(defclass linear-module object w)
(defmethod linear-module linear-module (ninputs noutputs)
  (setq w (matrix noutputs ninputs)))
(defmethod linear-module fprop (input output)
  (==> output resize (idx-dim :w:x 0))
  (idx-m2dotm1 :w:x :input:x :output:x) ())
(defmethod linear-module bprop (input output)
  (idx-m2dotm1 (transpose :w:x) :output:dx :input:dx)
  (idx-m1extm1 :output:dx :input:x :w:dx) ())
```



# Sigmoid Module (tanh: hyperbolic tangent)




- fprop:  $(X_{out})_i = \tanh((X_{in})_i + B_i)$
- bprop to input:  

$$\left(\frac{\partial E}{\partial X_{in}}\right)_i = \left(\frac{\partial E}{\partial X_{out}}\right)_i \tanh'((X_{in})_i + B_i)$$
- bprop to bias:  

$$\frac{\partial E}{\partial B_i} = \left(\frac{\partial E}{\partial X_{out}}\right)_i \tanh'((X_{in})_i + B_i)$$
- $\tanh(x) = \frac{2}{1+\exp(-x)} - 1 = \frac{1-\exp(-x)}{1+\exp(-x)}$

```
(defclass tanh-module object bias)
(defmethod tanh-module tanh-module 1
  (setq bias (apply matrix 1)))
(defmethod tanh-module fprop (input output)
  (==> output resize (idx-dim :bias:x 0))
  (idx-add :input:x :bias:x :output:x)
  (idx-tanh :output:x :output:x))
(defmethod tanh-module bprop (input output)
  (idx-dtanh (idx-add :input:x :bias:x) :input:dx)
  (idx-mul :input:dx :output:dx :input:dx)
  (idx-copy :input:dx :bias:dx) ()))
```

# Euclidean Module



- fprop:  $X_{out} = \frac{1}{2} \|X_{in} - Y\|^2$
- bprop to  $X$  input:  $\frac{\partial E}{\partial X_{in}} = X_{in} - Y$
- bprop to  $Y$  input:  $\frac{\partial E}{\partial Y} = Y - X_{in}$


```
(defclass euclidean-module object)
(defmethod euclidean-module run (input1 input2 output)
  (idx-copy :input1:x :input2:x)
  (:output:x 0) ())
(defmethod euclidean-module fprop (input1 input2 output)
  (idx-sqrdist :input1:x :input2:x :output:x)
  (:output:x (* 0.5 (:output:x))) ())
(defmethod euclidean-module bprop (input1 input2 output)
  (idx-sub :input1:x :input2:x :input1:dx)
  (idx-dotm0 :input1:dx :output:dx :input1:dx)
  (idx-minus :input1:dx :input2:dx))
```

# Assembling the Network: A single layer

---

```
;; One layer of a neural net
(defclass nn-layer object
  linear ; linear module
  sum ; weighted sums
  sigmoid ; tanh-module
)
(defmethod nn-layer nn-layer (ninputs noutputs)
  (setq linear (new linear-module ninputs noutputs))
  (setq sum (new state noutputs))
  (setq sigmoid (new tanh-module noutputs)) ())
(defmethod nn-layer fprop (input output)
  (=> linear fprop input sum)
  (=> sigmoid fprop sum output) ())
(defmethod nn-layer bprop (input output)
  (=> sigmoid bprop sum output)
  (=> linear bprop input sum) ())
```

# Assembling a 2-layer Net



- Class implementation for a 2 layer, feed forward neural net.

```
(defclass nn-2layer object
  layer1 ; first layer module
  hidden ; hidden state
  layer2 ; second layer
)

(defmethod nn-2layer nn-2layer (ninputs nhidden)
  (setq layer1 (new nn-layer ninputs nhidden))
  (setq hidden (new state nhidden))
  (setq layer2 (new nn-layer nhidden noutput))
)
```

# Assembling the Network: fprop and bprop

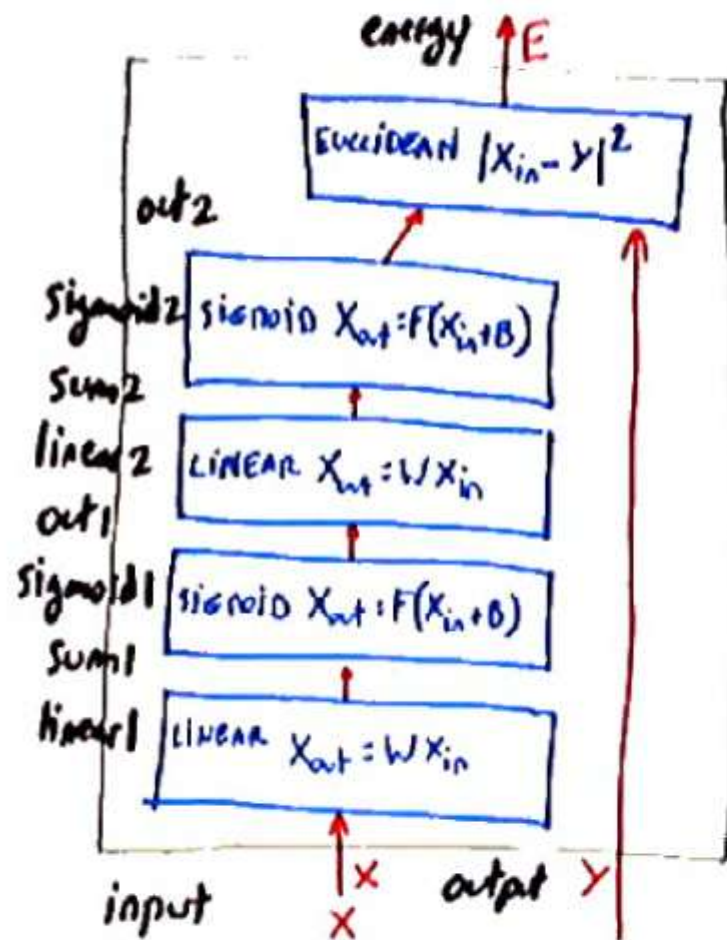
---

Implementation of a 2 layer, feed forward neural net.

```
(defmethod nn-2layer fprop (input output)
  (==> layer1 fprop input hidden)
  (==> layer2 fprop hidden output) ())
```


```
(defmethod nn-2layer bprop (input output)
  (==> layer2 bprop hidden output)
  (==> layer1 bprop input hidden) ())
```

# Assembling the Network: training



- A training cycle:
- pick a sample  $(X^i, Y^i)$  from the training set.
- call fprop with  $(X^i, Y^i)$  and record the error
- call bprop with  $(X^i, Y^i)$
- update all the weights using the gradients obtained above.
- with the implementation above, we would have to go through each and every module to update all the weights. In the future, we will see how to “pool” all the weights and other free parameters in a single vector so they can all be updated at once.

# Any Architecture works



- **Any connection is permissible**

- ▶ Networks with loops must be “unfolded in time”.

- **Any module is permissible**

- ▶ As long as it is continuous and differentiable almost everywhere with respect to the parameters, and with respect to non-terminal inputs.