Detecting Phase Transitions with Artificial Neural Networks

Sebastian J. Wetzel

Institute for Theoretical Physics, University of Heidelberg

2.5.2017, Cold Quantum Coffee, ITP Heidelberg

H G S F P

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Outline

- Invitation: Phase transitions from microscopic physics
- Method: Artificial neural networks
- Testing ground: Ising Model
- Results

Unsupervised learning of phase transitions: from principal component analysis to variational autoencoders
S. J. Wetzel ' 2017

Invitation: Phase transitions from microscopic physics

Hamiltonian

$$H(S) = -J \sum_{<ij>_{NN}} s_i s_j$$

Ising Model

Order Parameter

$$M(S) = \frac{1}{N} \sum_i s_i$$

$$\bar{M}(T) = \frac{1}{Z} \sum_{S \in \Lambda} M(S) \exp(-H(S)/T)$$

Goal:

- Phase Diagram

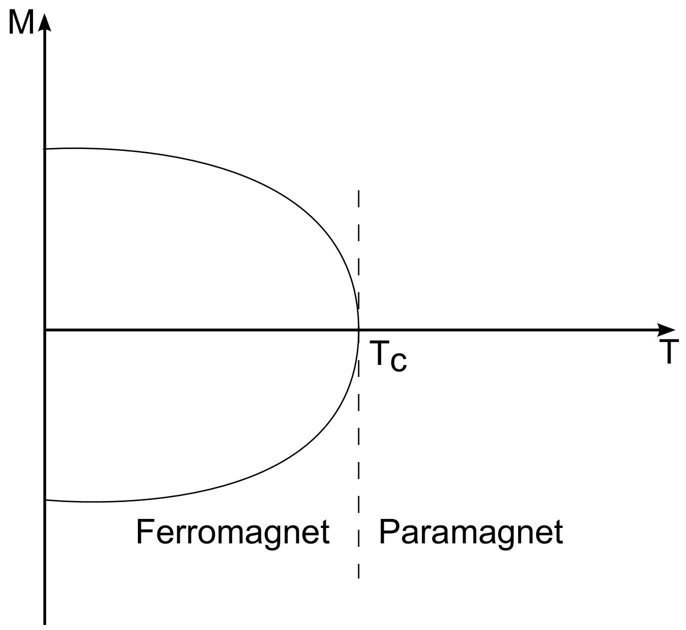

Invitation: Phase transitions from microscopic physics

Hamiltonian

$$H(S) = -J \sum_{<ij>_{NN}} s_i s_j$$

Ising Model

Order Parameter

$$M(S) = \frac{1}{N} \sum_i s_i$$

$$\bar{M}(T) = \frac{1}{Z} \sum_{S \in \Lambda} M(S) \exp(-H(S)/T)$$

Monte Carlo Sampling

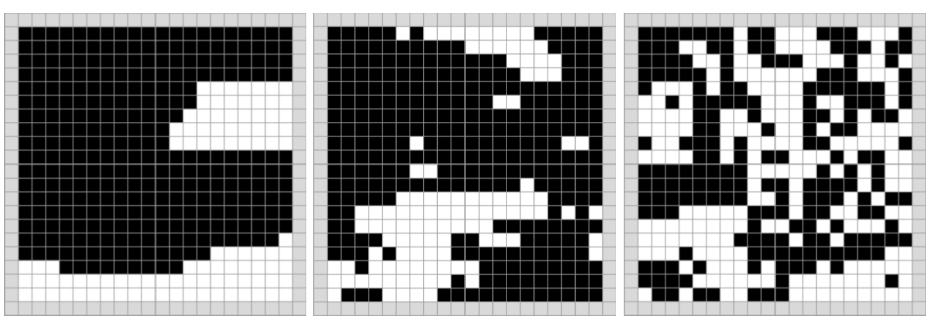

Wetterich Equation

$$\partial_k \Gamma_k = \frac{1}{2} \mathrm{Tr} \left(\Gamma^{(2)} + R_k\right)^{-1} \partial_k R_k$$

Invitation: Phase transitions from microscopic physics

Hamiltonian

$$H(S) = -J \sum_{<ij>_{NN}} s_i s_j$$

Order Parameter

$$M(S) = \ldots$$

$$\bar{M}(T) = \sum_{S\in\Lambda} \ldots) \exp(-H(S)/T)$$

Unknown?
Hard to find?
Hard to define?

Monte Carlo Sampling

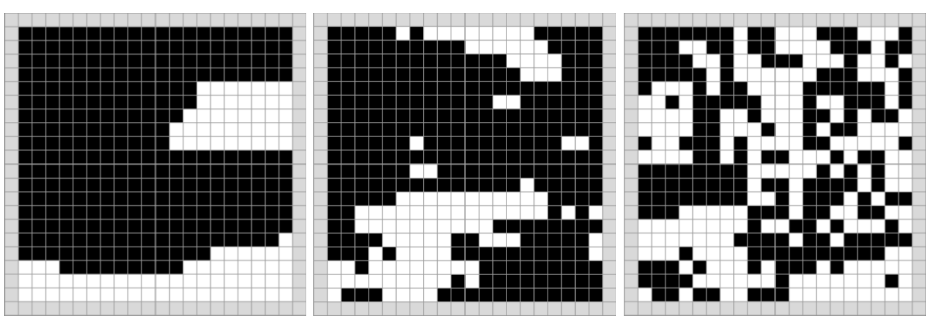

Wetterich equation

$$\partial_k \Gamma_k = \frac{1}{2} \mathrm{Tr}\,(\Gamma^{(2)} + R_k)^{-1} \partial_k R_k$$

Invitation: Phase transitions from microscopic physics

Hamiltonian

$$H(S) \ldots \sum_{\ldots >_{NN}} s_i s_j$$

Experiment?
Hamiltonian
unknown?

Order Parameter

$$M(S) = \ldots$$

$$\bar{M}(T) = \ldots \sum_{S \in \Lambda} \ldots \exp(-H(S)/T)$$

Unknown?
Hard to find?
Hard to define?

Monte Carlo Sampling

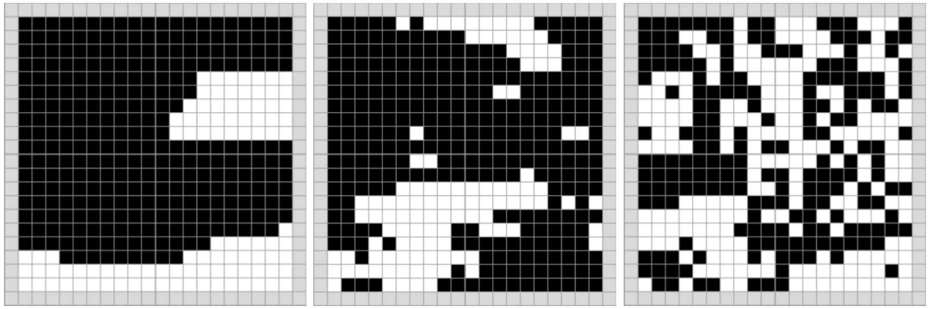

Wetterich equation

$$\partial_k \Gamma_k = \frac{1}{2} \mathrm{Tr} \left(\Gamma^{(2)} + R_k \right)^{-1} \partial_k R_k$$

Invitation: Phase transitions from microscopic physics

Hamiltonian

$H(S) \ldots \sum_{\ldots >_{NN}} s_i s_j$

Experiment?
Hamiltonian unknown?

Order Parameter

$M(S) = \ldots$

$\bar{M}(T) = \ldots \sum_{S \in \Lambda} \ldots \exp(-H(S)/T)$

Unknown?
Hard to find?
Hard to define?

Monte Carlo Sampling

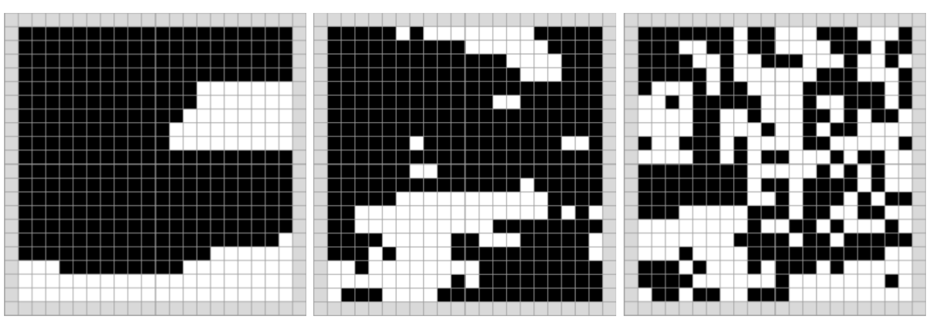

Wetterich equation

$$\partial_k \Gamma_k = \frac{1}{2} \mathrm{Tr} \left(\Gamma^{(2)} + R_k\right)^{-1} \partial_k R_k$$

Possible? Solution: use Artificial Neural Networks!

Machine Learning

„Machine learning is the subfield of computer science that gives computers the ability to learn without being explicitly programmed.“ - Wikipedia

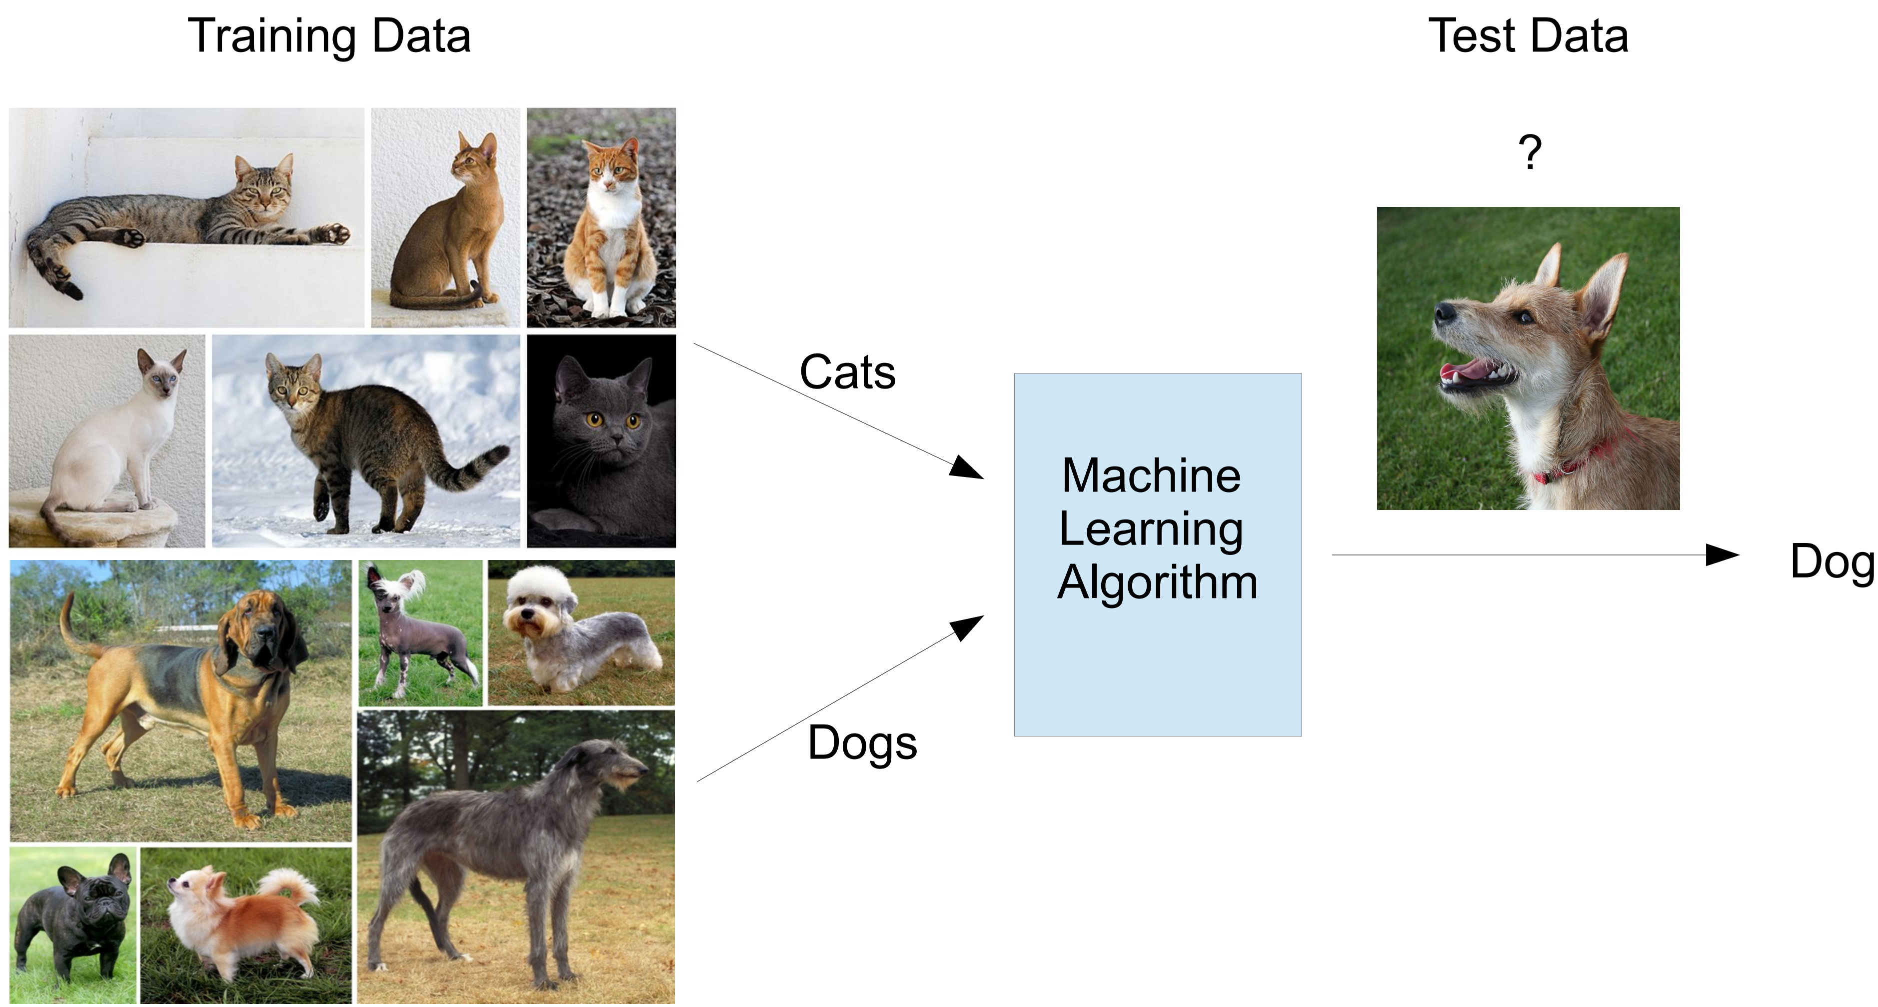

Artificial Neural Networks

Feed forward neural network

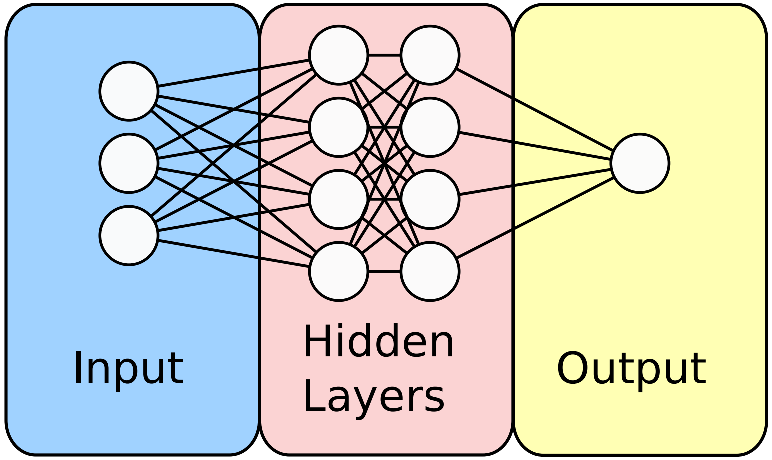

Input: Data $X = (x_1, ..., x_d)$, Label Y

Output: $Y_{pred} = F(X)$

Goal: find F such that $Y_{pred} \approx Y$

Artificial Neural Networks

Perceptron

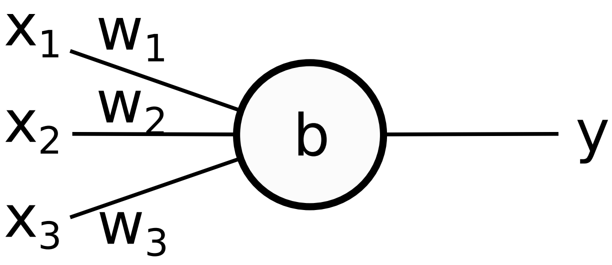

$$y = f(\vec{x} \cdot \vec{w} + b)$$

Example: Buying a house

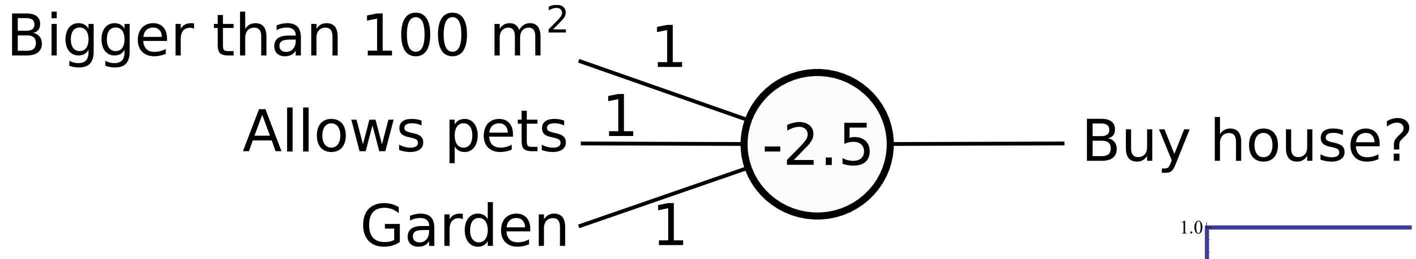

$$y = \Theta(\vec{x} \cdot \vec{w} + b)$$

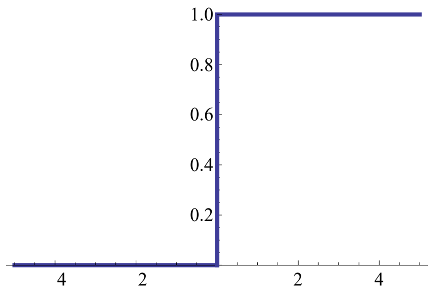

- If all 3 conditions are fulfilled the perceptron decides to buy

Artificial Neural Networks

Activation functions in neural networks

Rectified linear unit (relu)	$\max(0, x)$	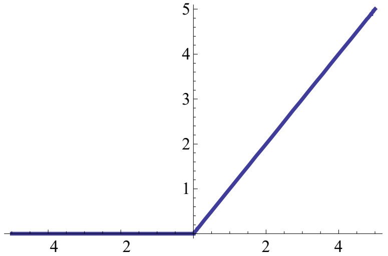	Common interlayer activation function
Sigmoid	$\frac{1}{1 + \exp(-x)}$	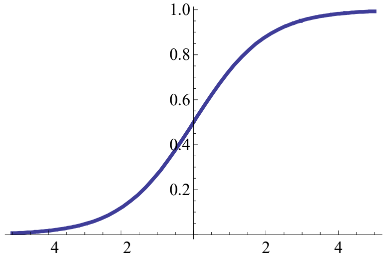	Predicting probabilities of discrete variables
tanh	$\tanh(x)$	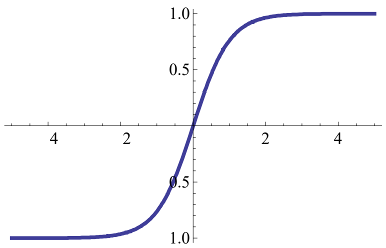	Predicting an output constrained by an interval

Training

Objective functions (loss functions)

- Eg mean squared error (average over all samples)

$$MSE = \frac{1}{N}\sum_i \left(Y_{(i)} - F(X_{(i)})\right)^2$$

Training

- Determination of w_{ij}^L and b_i^L
- Gradient descent

$$\frac{\partial MSE}{\partial w_{ij}^L} \text{ and } \frac{\partial MSE}{\partial b_i^L}$$

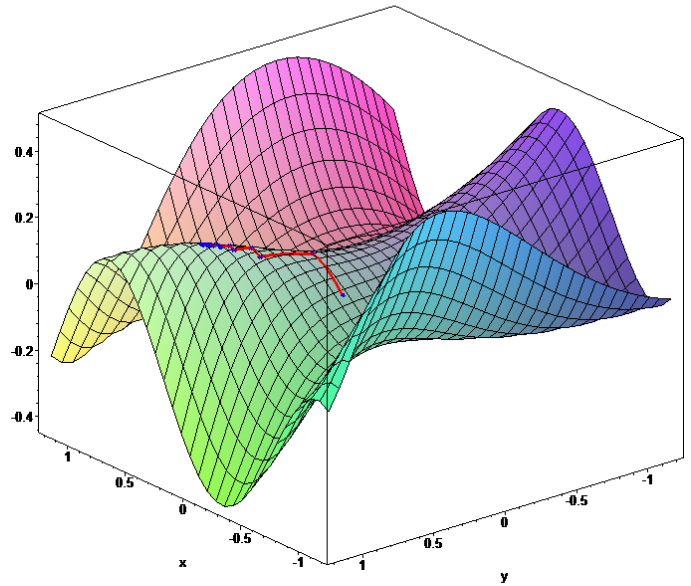

- Backpropagation algorithm

Supervised Learning of Phase Transitions
2d Ising Model

- Data: Monte Carlo samples
- Training at well known points in phase diagram
- Labels: Phase
- Testing in interval containing phase transition
- Estimate within 1% of exact value $T_c = \frac{2}{\ln(1+\sqrt{2})}$

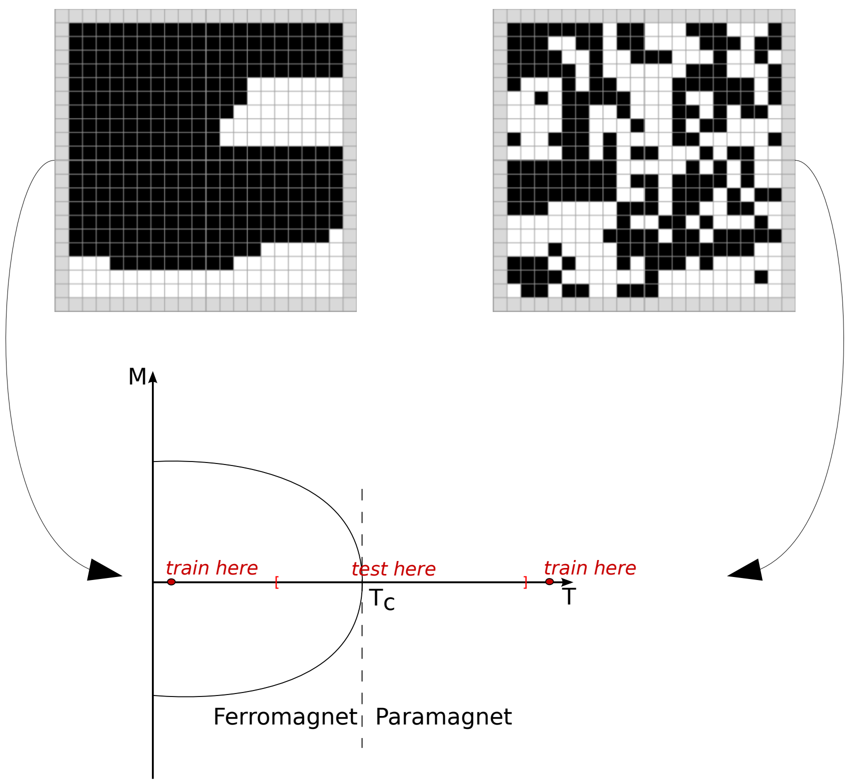

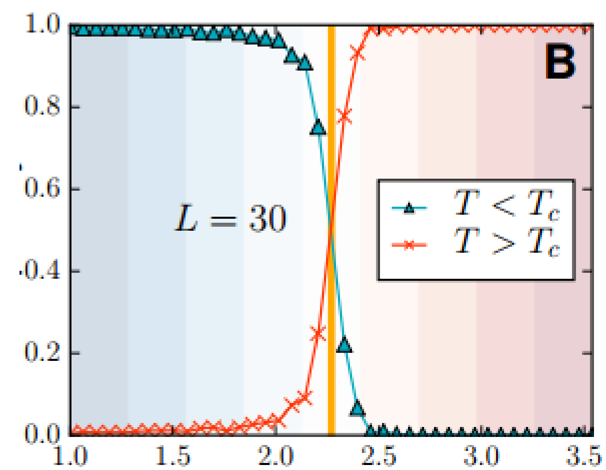

Machine Learning Phases of Matter
Carrasquilla, Melko ' 2016

Supervised Learning of Phase Transitions

Limitations of Supervised Learning

- Example Hubbard Model:
  rich phase diagram,
  many unknown phases
  - Pseudogap?
  - Strange Metal?
  - Coexistence of AF and SC?
- Detecting unknown phases?
- In order to determine the phase transition, you already need to know the existence

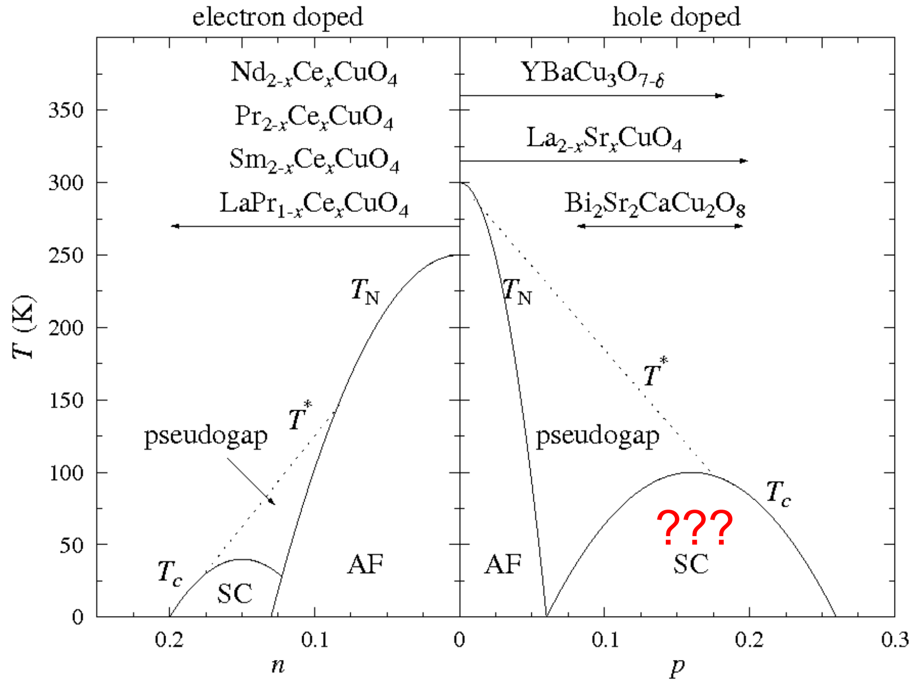

Unsupervised Learning

Up to now we discussed supervised learning,
where labels were given for training.
Now we transition to unsupervised Learning.

Training Data

?

?

Machine Learning Algorithm

Test Data

?

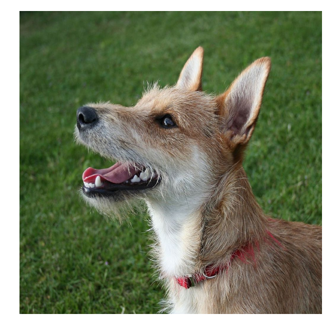

???

Unsupervised Learning

Up to now we discussed supervised learning,
where labels were given for training.
Now we transition to unsupervised Learning.

Training Data

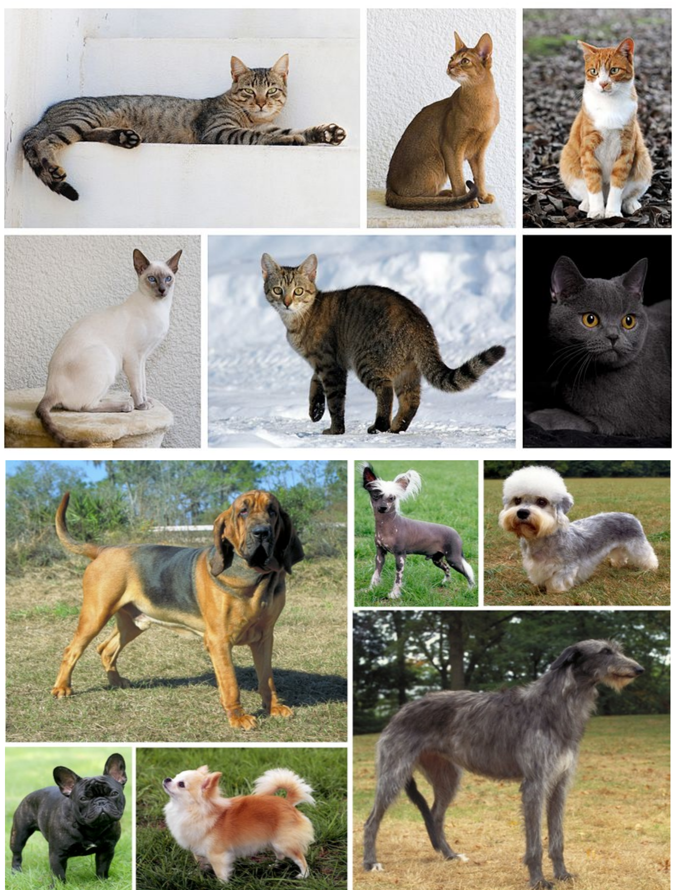

?

?

Machine Learning Algorithm

Clustering of Dog and Cat Images

Test Data

?

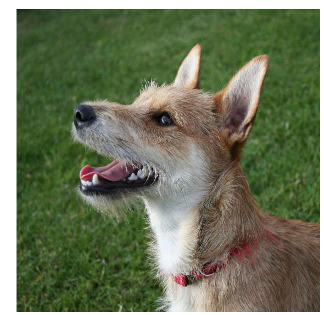

Cluster 2

Cluster 1: Cats
Cluster 2: Dogs

Unsupervised Learning of Phase transitions

Method	Invented	Phase transitions
Principal component analysis	*K. Pearson 1901*	*L. Wang 2016*
Kernel Principal component analysis	*Schollkopf, Smola, Müller 1999*	
Autoencoder	*LeCun 1987 , Bourlard, Kamp 1988*	*S.J. Wetzel 2017*
Variational Autoencoder	*Kingma, Welling 2013*	

+Non-Linear Features

+Scaling to huge Datasets
-Overfitting

+Less Overfitting
+Latent Parameter Model

Autoencoder

- Architecture: Encoder NN + Decoder NN
- Objective: Minimize Reconstruction error
- Bottleneck: Latent Variables

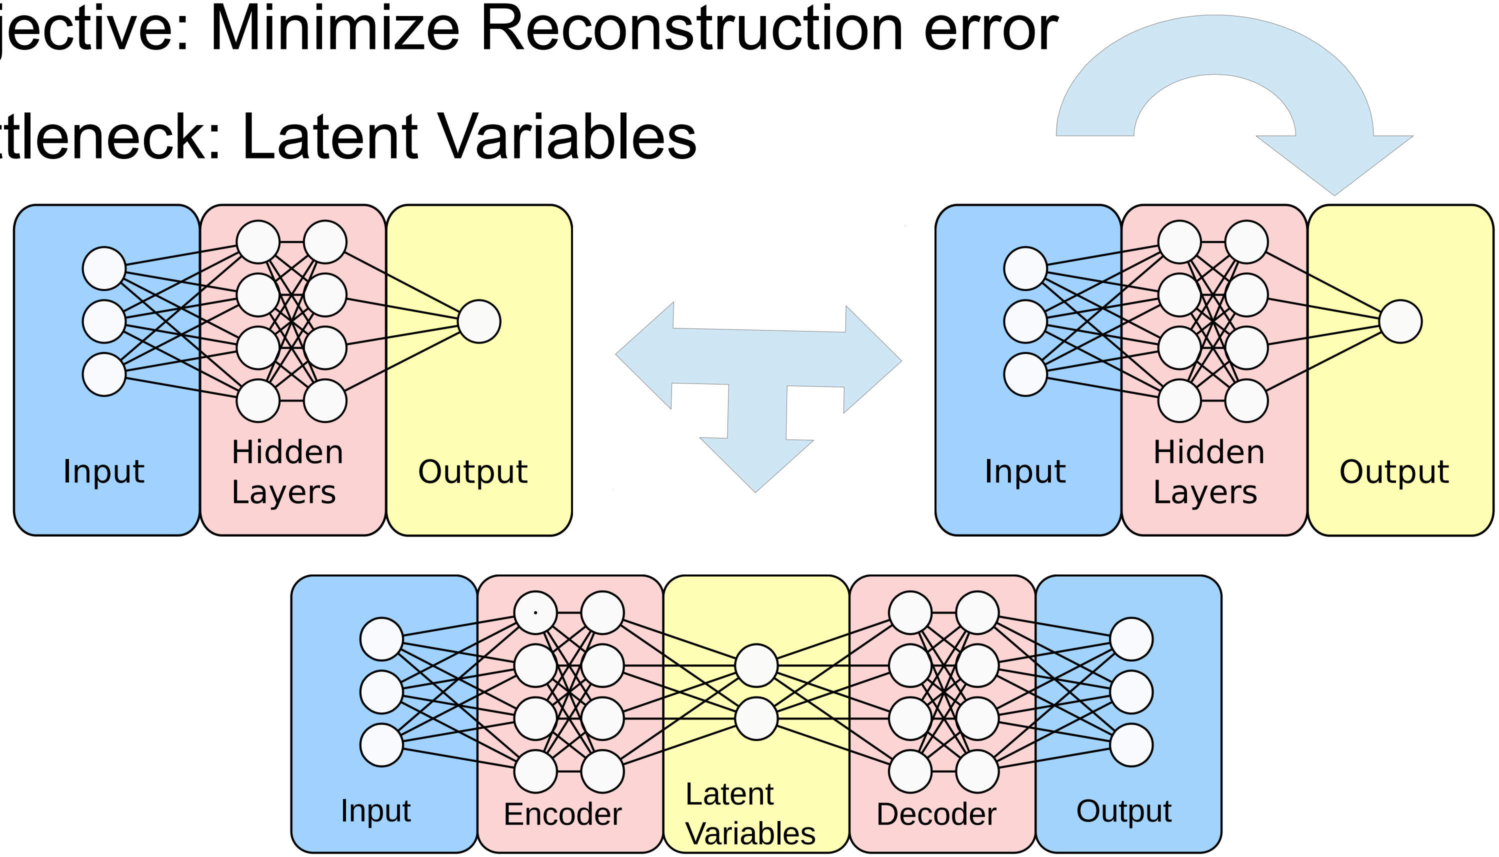

$$MSE = \frac{1}{N} \sum_i \left\| Y_{(i)} - F(X_{(i)}) \right\|^2$$

What do Autoencoders store?
2d Ising Model

- Interesting quantities:
  - Reconstructions of the samples
  - Physical interpretation of the latent parameters

  Correlation between latent parameter and the magnetization

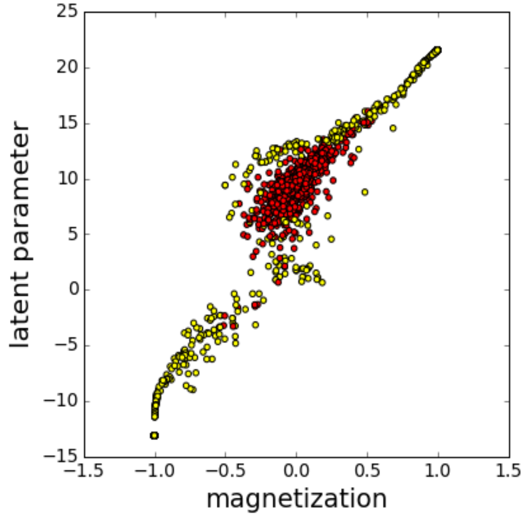

- Problems:
  - Very hard to infer order parameter from this diagram
  - Latent parameter can in principle store many substructures seen on the data

Variational Autoencoder

- Architecture: Encoder NN + Decoder NN
- Assumes data can be generated from Gaussian prior
- Input X is encoded into latent variables Z which are decoded producing the output Y

$$MSE = \frac{1}{N} \sum_i \left\| Y_{(i)} - F(X_{(i)}) \right\|^2$$

$$KLO = D_{KL}(\mathcal{N}(mean(Z), var(Z)) \| \mathcal{N}(0,1))$$

- Can be understood as a regularization of the traditional autoencoder
- Training makes sure that neighboring latent representations encode similar configurations

From Autoencoders to Variational Autoencoders

- Why do we need a variational autoencoder?
  - We approximate 1 to 1 mapping to the order parameter

AE:

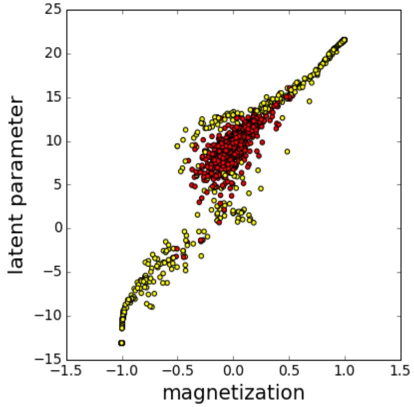

VAE:

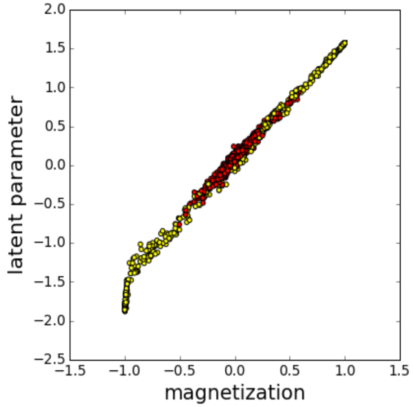

How to determine an optimal number of latent neurons

- No theory
- Try different numbers
- Look for small ranges

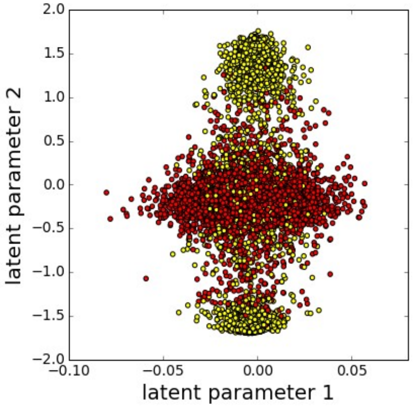

Variational Autoencoder

Why could this work?

- Autoencoder encodes the general structure of samples in the decoder
- The latent variables store the parameters that hold the most information about quantifiable structures on configurations
- In the unordered phase sample configurations differ by random entropy fluctuations. The variational autoencoder averages over these fluctuations and thus fails to learn a quantity which quantifies these structures
- In the ordered phase the variational autoencoder learns a common correlation between the spins, whose strength is quantified by a latent variable with coincides with the order parameter

Variational Autoencoder

Why could this work?

- Autoencoder encodes the general structure of samples in the decoder
- The latent variables store the parameters that hold the most information about quantifiable structures on a configurations
- In the unordered phase sample configurations differ by random entropy fluctuations. The variational autoencoder averages over these fluctuations and thus fails to learn a quantity which quantifies these structures
- In the ordered phase the variational autoencoder learns a common correlation between the spins, whose strength is quantified by a latent variable with coincides with the order parameter
- Reconstruction Error as Universal Identifier for Phase Transitions

Results
2d Ising Model

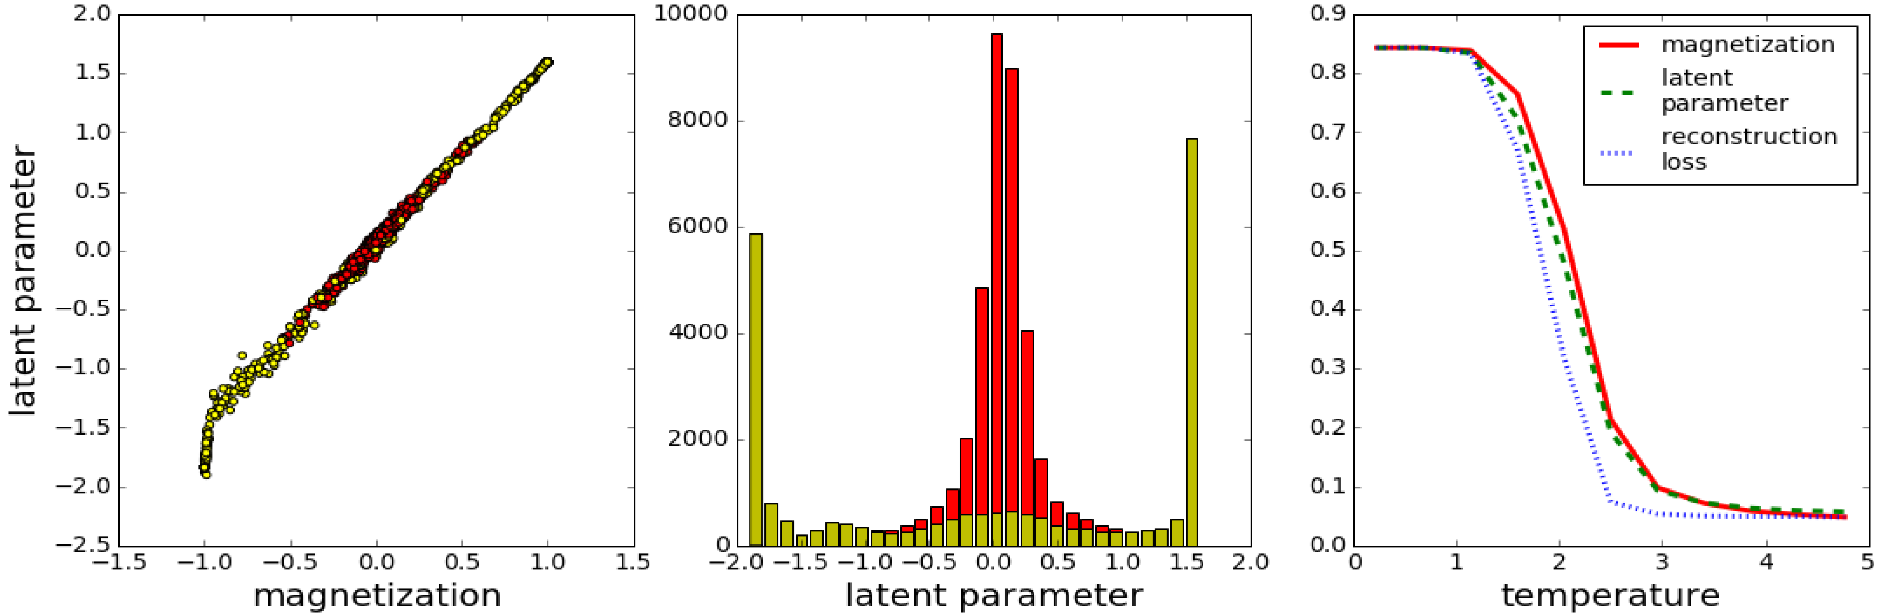

Ferromagnetic Ising model on the square lattice

- Latent parameter corresponds to magnetization
- Identification of phases: Latent representations are clustered
- Location of phases: Magnetization, latent parameter and reconstruction loss show a steep change at the phase transition.

Results
2d Antiferromagnetic Ising Model

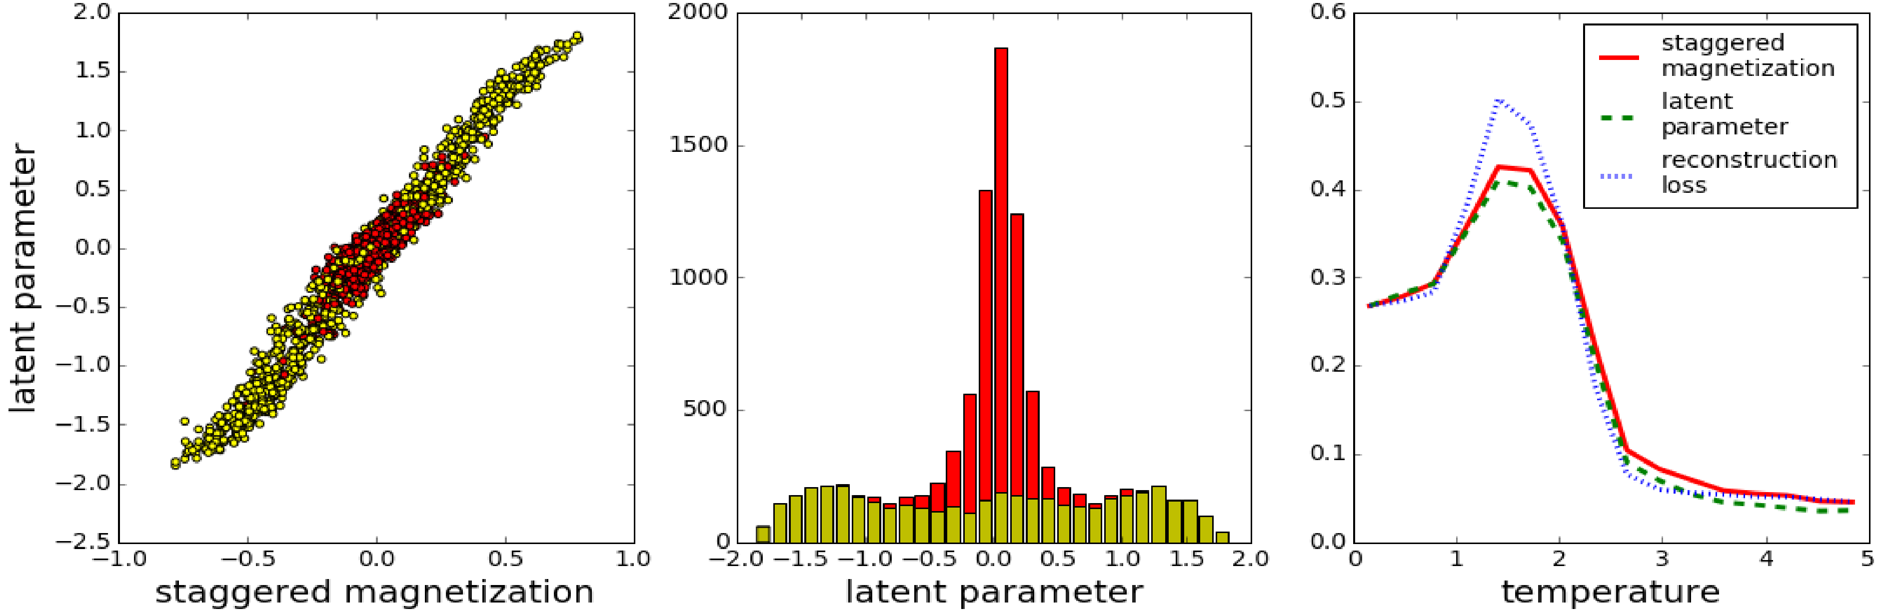

Antiferromagnetic Ising Model on the square lattice

- Spins correspond to order parameter depending on site
- Latent parameter corresponds to staggered magnetization
- Identification of phases: Latent representations are clustered
- Location of phases: Staggered magnetization, latent parameter and reconstruction loss show a steep change at the phase transition.

Results

3d XY Model

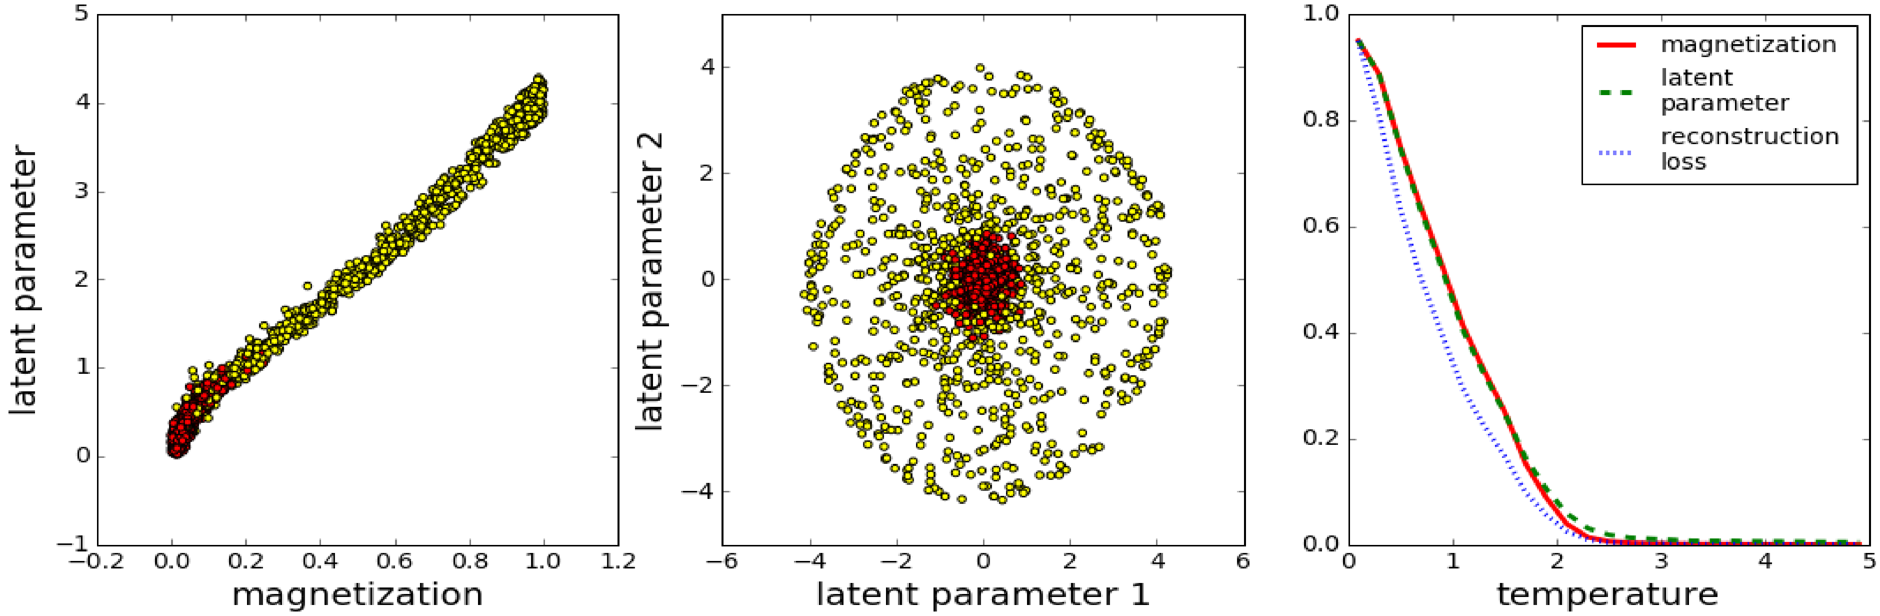

Ferromagnetic XY Model in 3d

- Continous phases have infinitely many representations
- Latent parameter corresponds to magnetization
- Identification of phases: Clustering could be inferred
- Location of phases: Magnetization, latent parameter and reconstruction loss show a steep change at the phase transition.

Conclusion

- Methods to pin down phase transitions, supervised learning
- Methods to detect phases, unsupervised learning
  - Latent parameter coincides with order parameter
  - Universal identifier: reconstruction error
- Caveat:
  - No proof
  - Requires huge amounts of sample configurations

Outlook

- More Complicated Systems
- Non-Local Order Parameters
- Interpretability of Order Parameters

Outlook

- More Complicated Systems
- Non-Local Order Parameters
- ~~Interpretability of Order Parameters~~
- Explicit expressions of Order Parameters

Machine Learning of Explicit Order Parameters at the Example of SU(2) Lattice Gauge Theory
S. J. Wetzel, M. Scherzer ' 2017 (in Preparation)