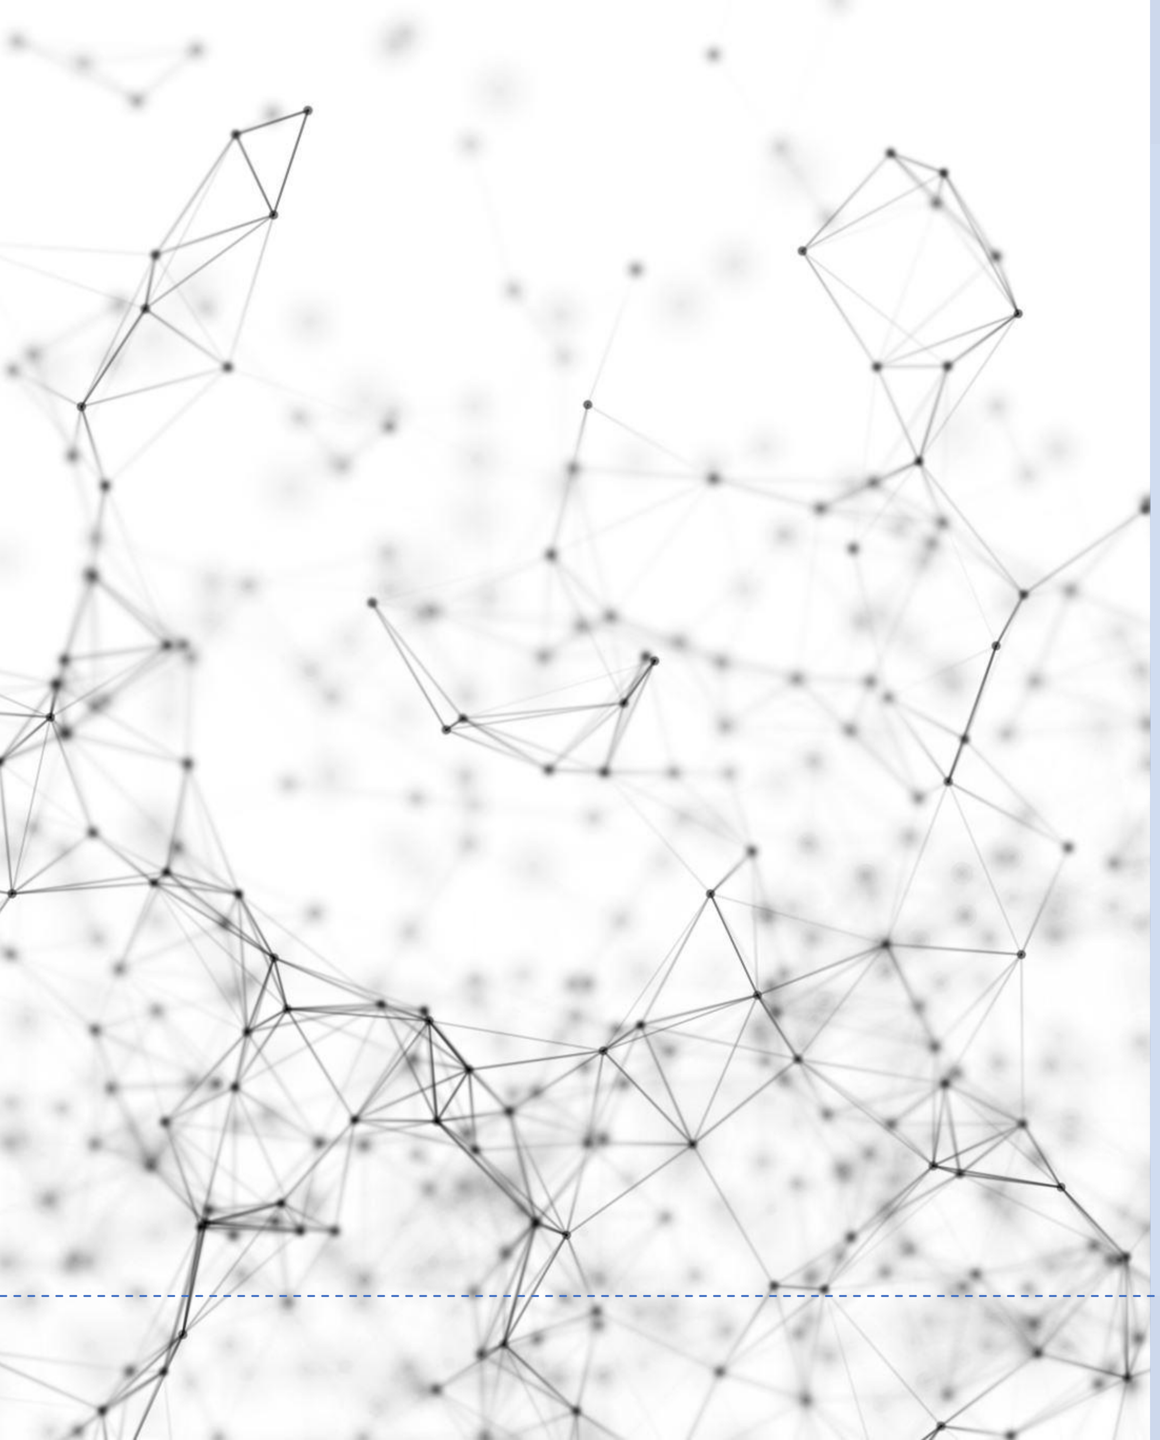

Learning Physical Concepts using Neural Networks

Jiangtian Yao, July 2021

This is based on…

PHYSICAL REVIEW LETTERS **124,** 010508 (2020)

Editors' Suggestion | Featured in Physics

Discovering Physical Concepts with Neural Networks

Raban Iten,[*,†] Tony Metger,[*,‡] Henrik Wilming, Lídia del Rio, and Renato Renner
ETH Zürich, Wolfgang-Pauli-Strasse 27, 8093 Zürich, Switzerland

(Received 17 July 2019; published 8 January 2020)

Despite the success of neural networks at solving concrete physics problems, their use as a general-purpose tool for scientific discovery is still in its infancy. Here, we approach this problem by modeling a neural network architecture after the human physical reasoning process, which has similarities to representation learning. This allows us to make progress towards the long-term goal of machine-assisted scientific discovery from experimental data without making prior assumptions about the system. We apply this method to toy examples and show that the network finds the physically relevant parameters, exploits conservation laws to make predictions, and can help to gain conceptual insights, e.g., Copernicus' conclusion that the solar system is heliocentric.

DOI: 10.1103/PhysRevLett.124.010508

Theoretical physics, like all fields of human activity, is influenced by the schools of thought prevalent at the time of development. As such, the physical theories we know may not necessarily be the simplest ones to explain experimental data, but rather the ones that most naturally followed from a relevant variables are or that dynamics should be described by differential equations. In certain situations one might not have such prior knowledge or does not want to impose it to allow the machine to find entirely different representations of the physical system.

Outline

- Introduction to Modeling
- Variational Auto Encoder (VAE) vs. SciNet
- Examples of Physical Problems
- Technical Aspects of the Optimization Procedure

Introduction

Time [s]	Position [m]
0	10
10	265
18	450
39	985
61	1550
…	…

Modelling: Observational Input

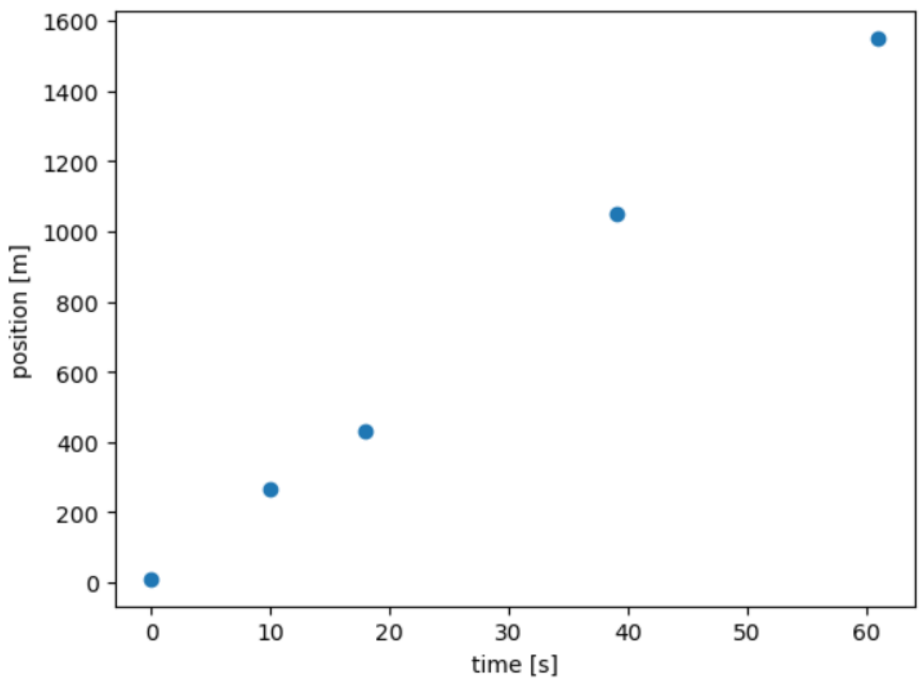

Time [s]	Position [m]
0	10
10	265
18	450
39	985
61	1550
…	…

Modelling: Postulate a Model

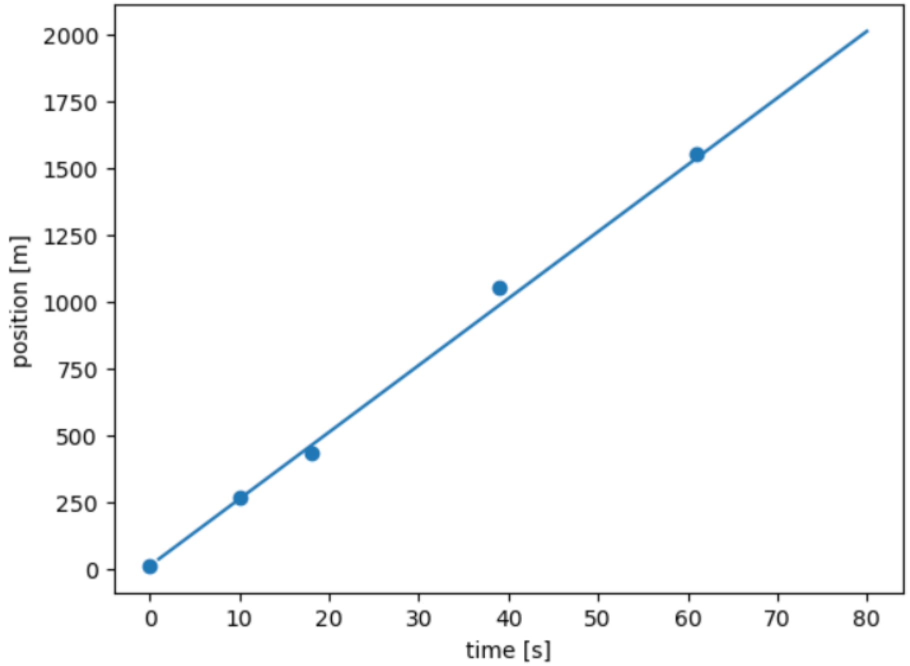

Time [s]	Position [m]
0	10
10	265
18	450
39	985
61	1550
...	...

Modelling: Extracting Parameters

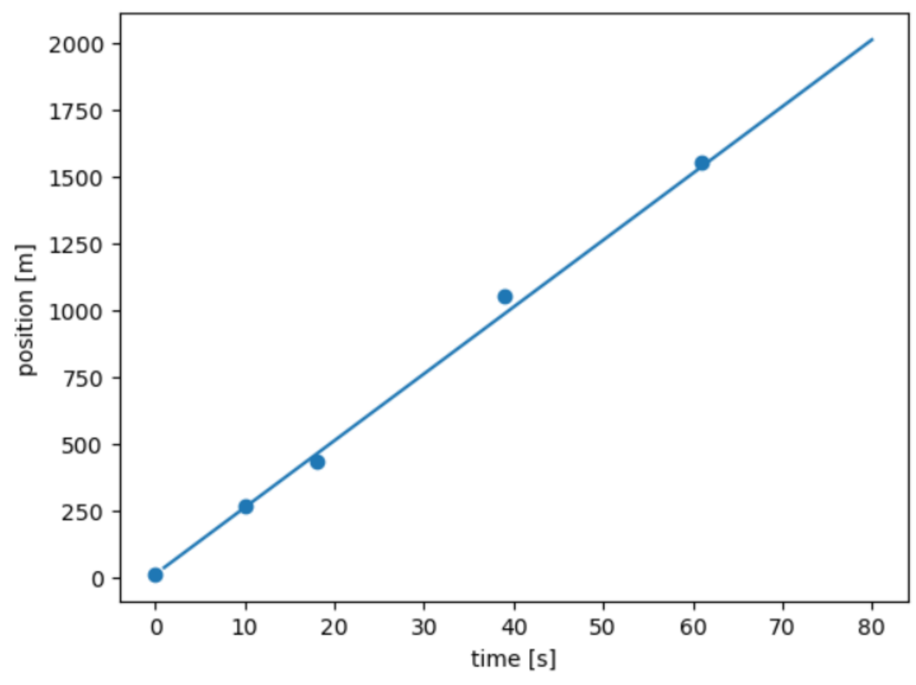

$x = v \cdot t + x_0$, $v = 25\, ms^{-1}, x_0 = 10\, m$

Model (Linear)

Parameter Values

Modelling: Extracting Parameters

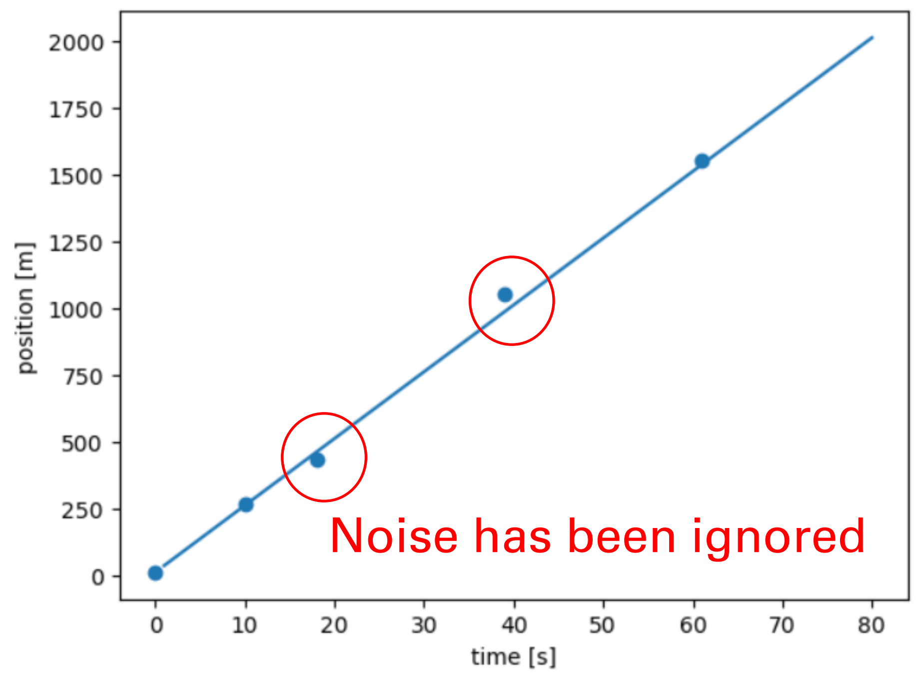

$x = v \cdot t + x_0$, $v = 25\, ms^{-1}, x_0 = 10\, m$

Model (Linear)

Parameter Values

Modelling: Making Predictions

Time [s]	Position [m]
100	?
200	?
…	…

Q: What is the car‘s position at time t=100 s ?

$x = v \cdot t + x_0$, $v = 25\, ms^{-1}, x_0 = 10\, m$

Model (Linear)

Parameter Values

Modelling: Making Predictions

Time [s]	Position [m]
100	2510
200	5010
…	…

A: 5010 m.

$$x = v \cdot t + x_0 \,, v = 25\, ms^{-1}, x_0 = 10\, m$$

Model (Linear)

Parameter Values

Modeling: Human -> Machine

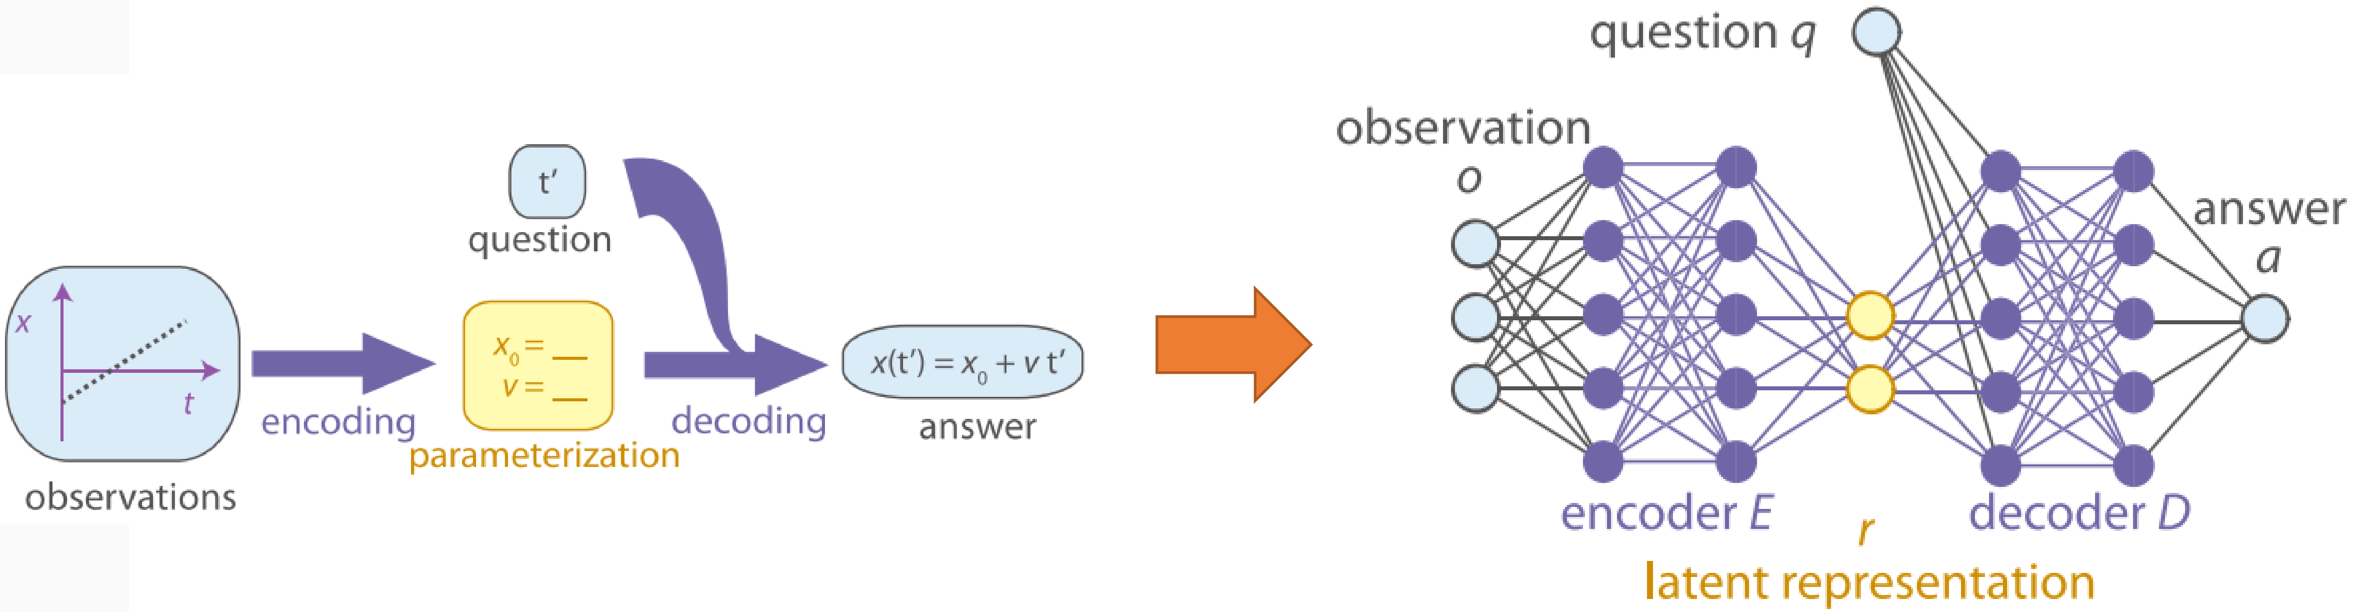

Modeling: Human -> Machine

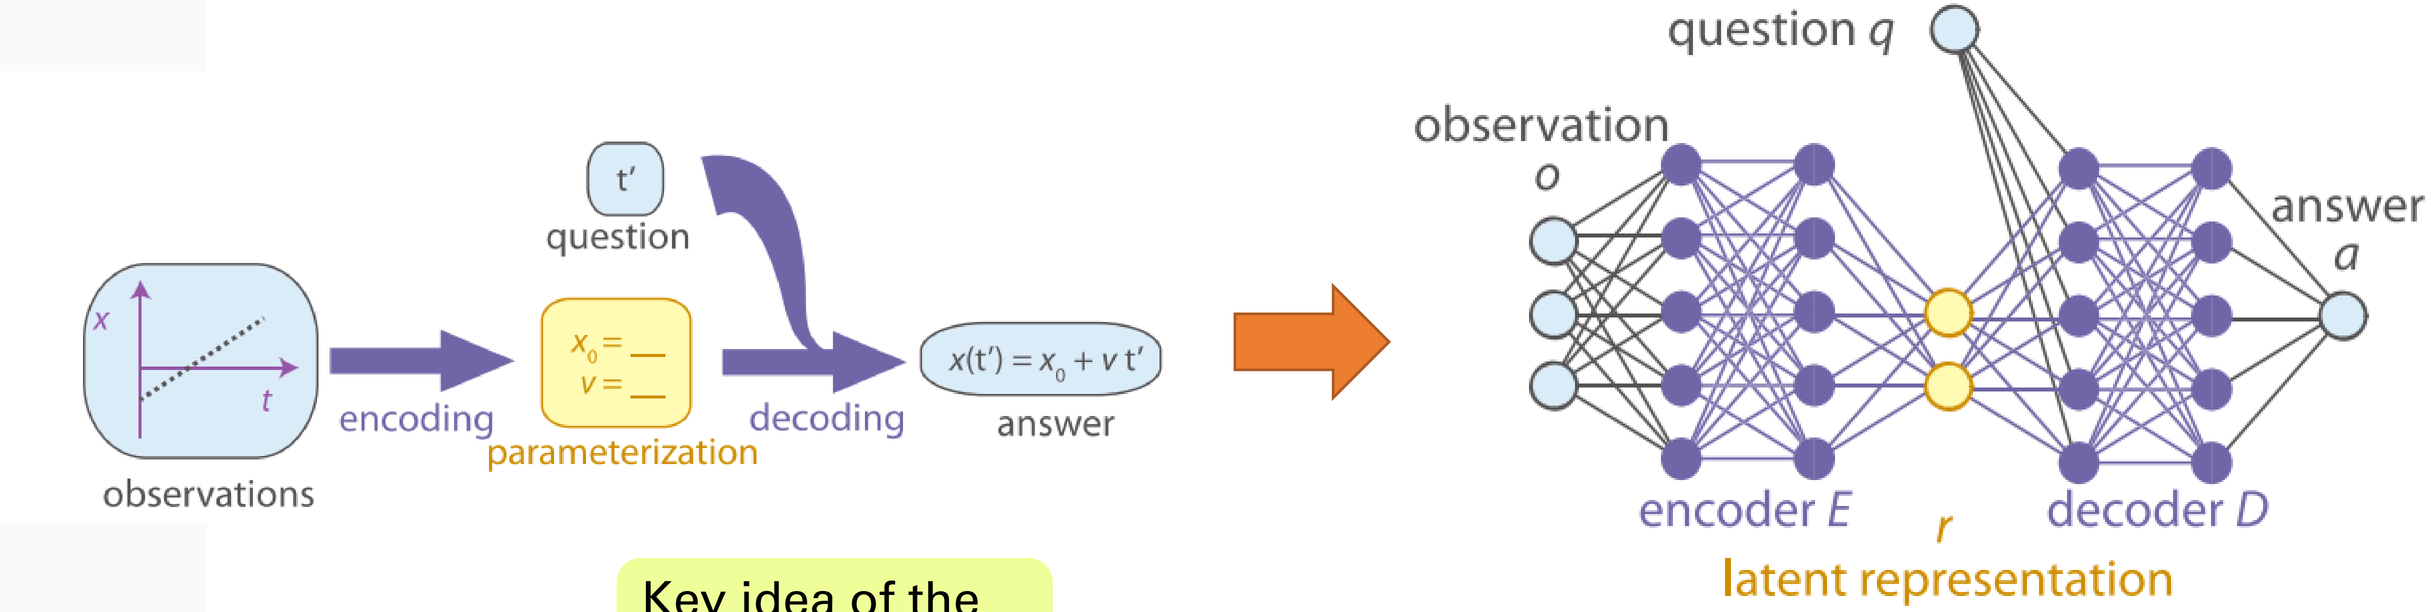

Key idea of the paper

Modeling: Human -> Machine

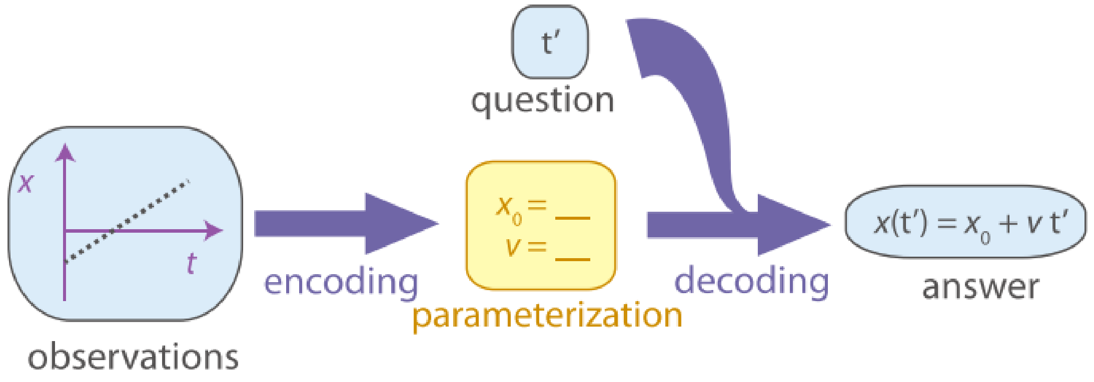

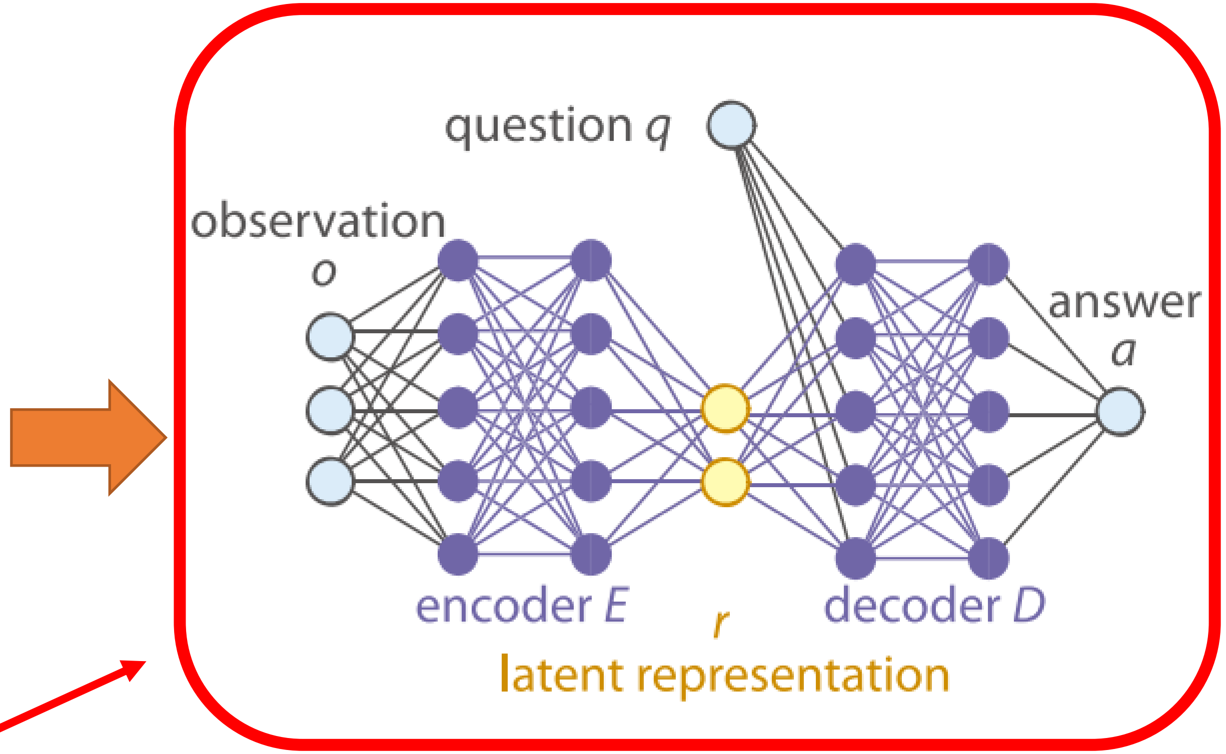

„SciNet“

Variational Auto Encoder (VAE)

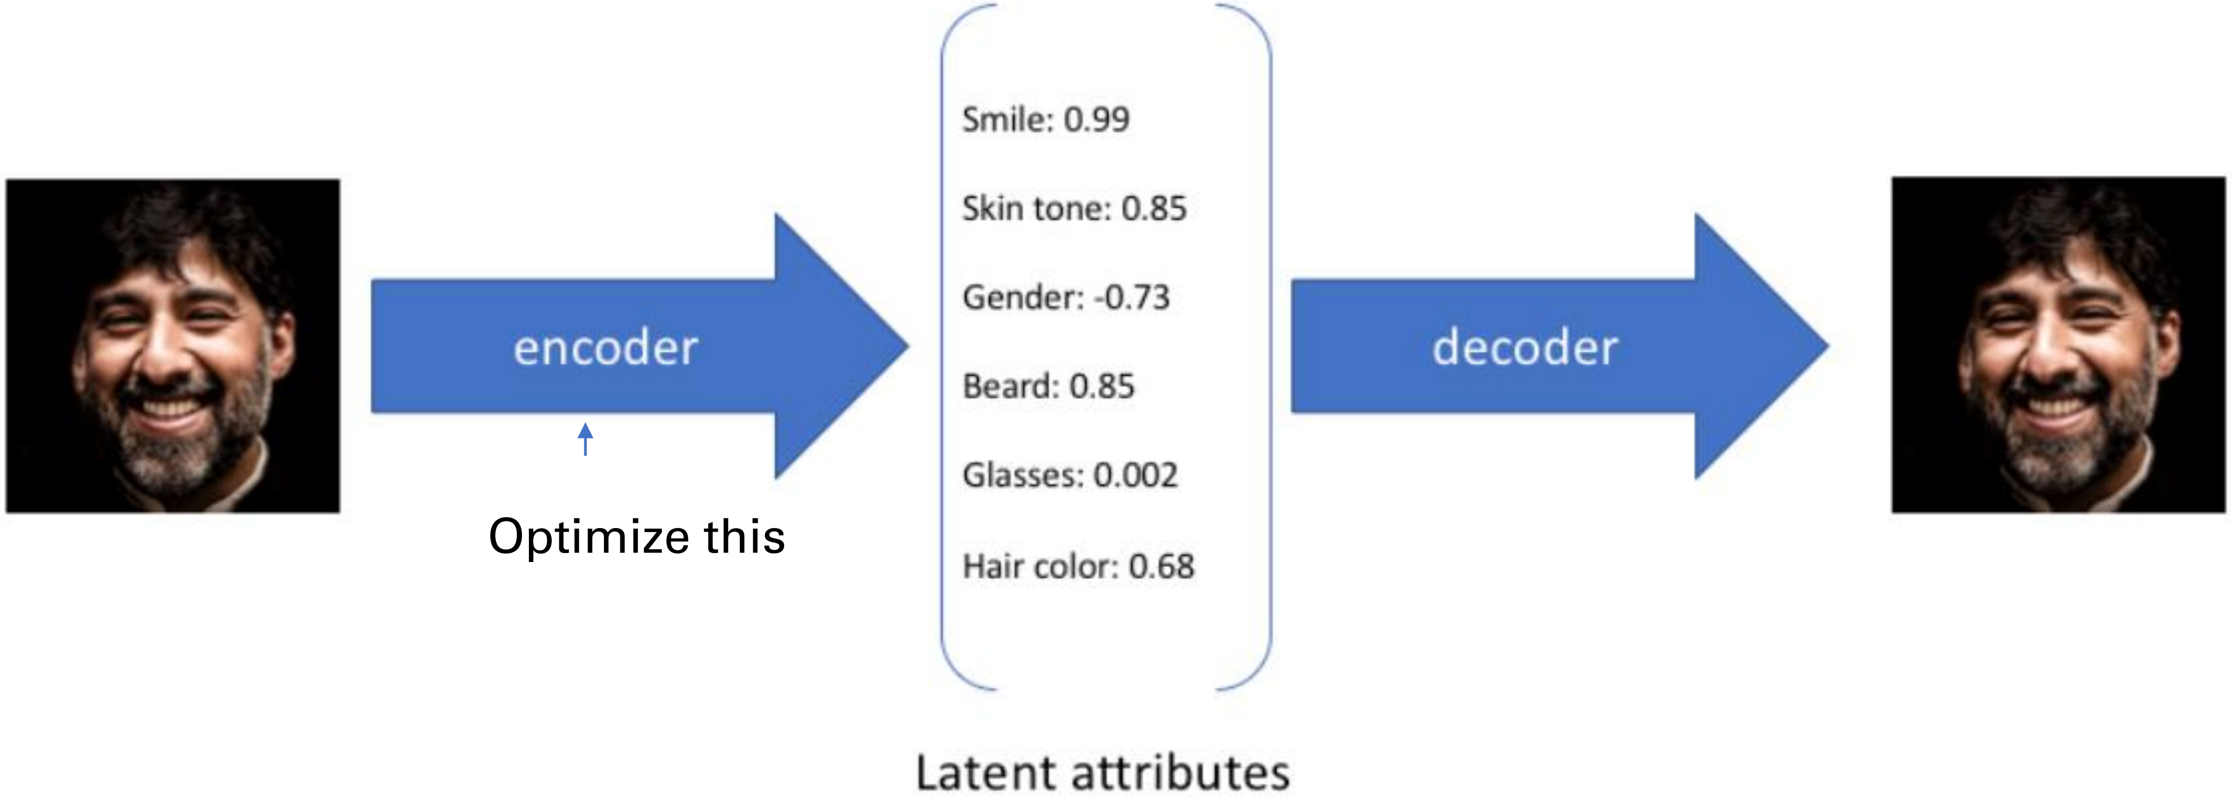

VAE Network Structure

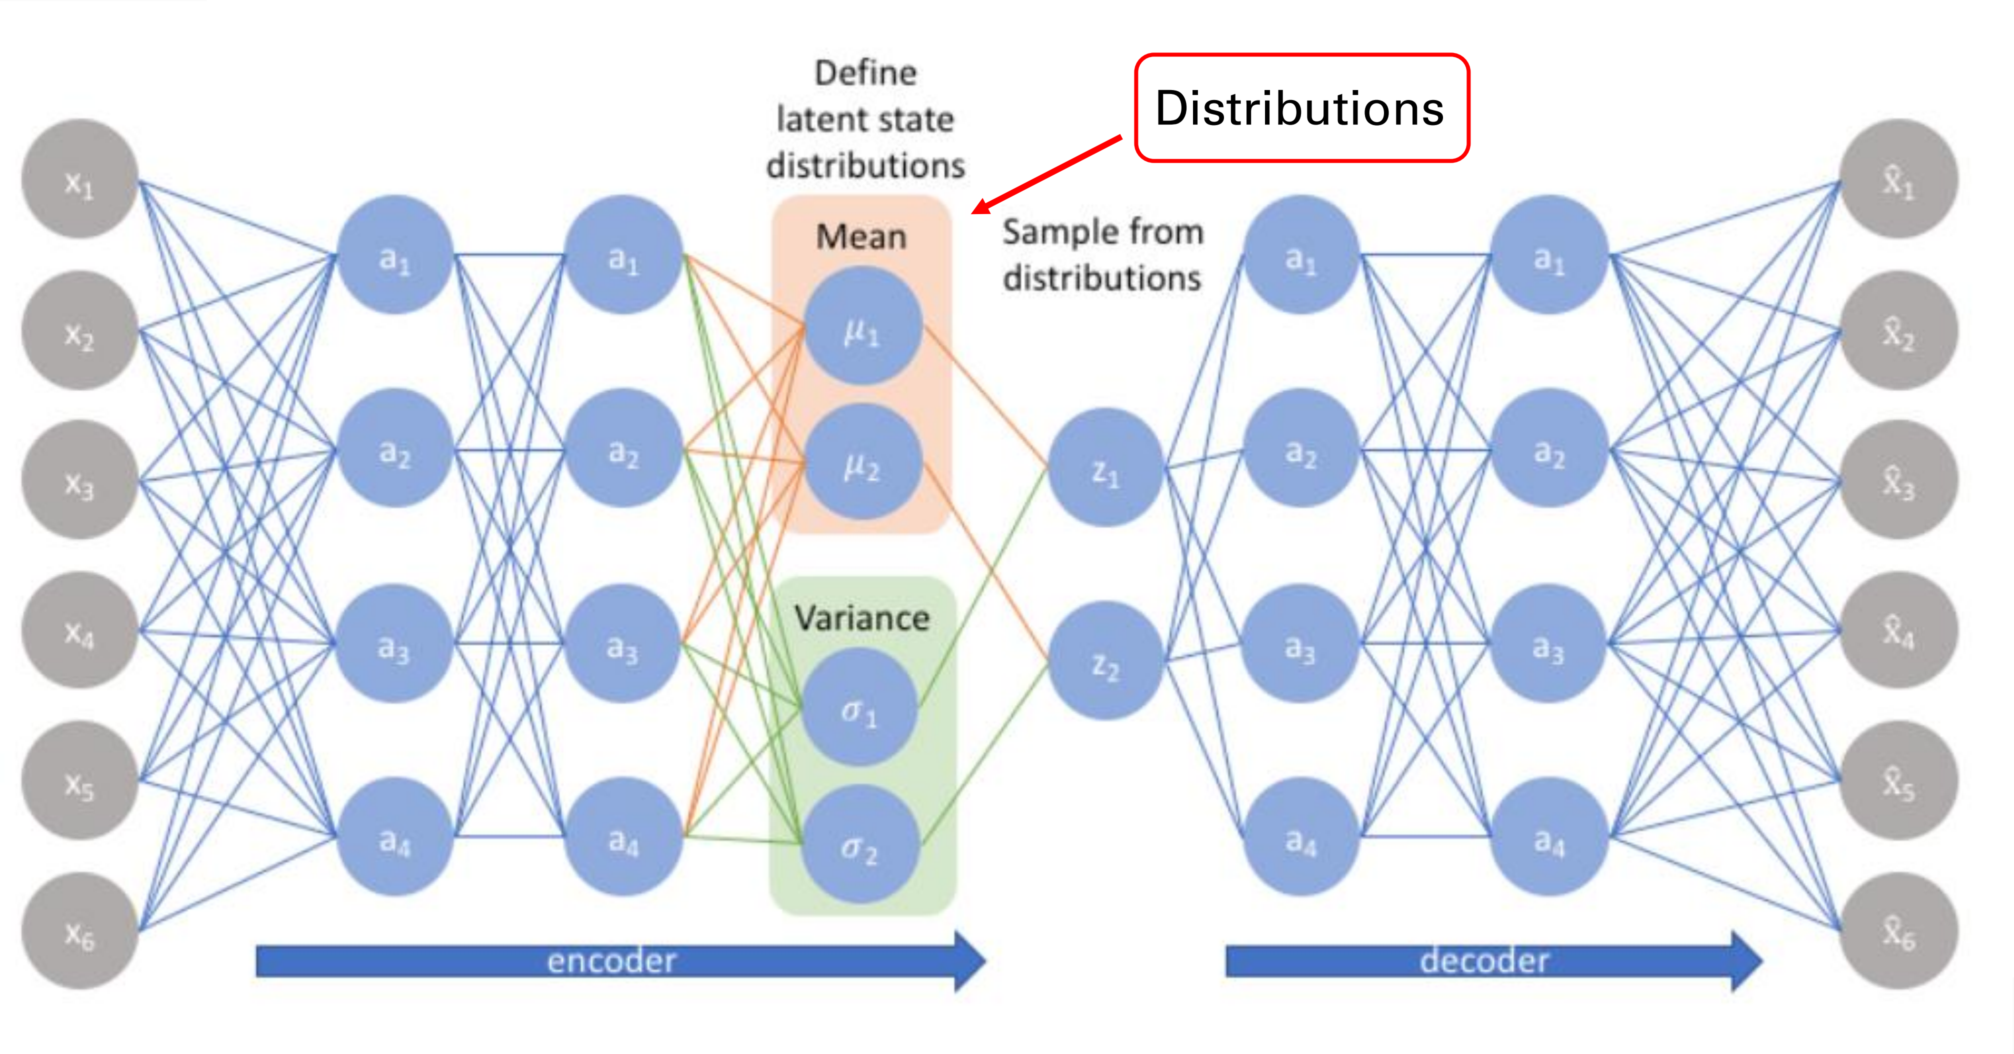

VAE: Criteria

- Maximize: faithfulness to prior distribution + reconstruction likelihood

VAE: Criteria

- Maximize: faithfulness to prior distribution + reconstruction likelihood

VAE: Criteria

- Maximize: faithfulness to prior distribution + reconstruction likelihood

$$E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\left(q_\phi(z|x) \middle|\middle| p(z)\right)$$

VAE: Criteria

- Maximize: faithfulness to prior distribution + reconstruction likelihood

$$E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\big(q_\phi(z|x) \| p(z)\big)$$

VAE: Criteria

- Maximize: faithfulness to prior distribution + reconstruction likelihood

$$E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\left(q_\phi(z|x)\middle|\middle|p(z)\right)$$

Will go into more detail later!

VAE vs. SciNet

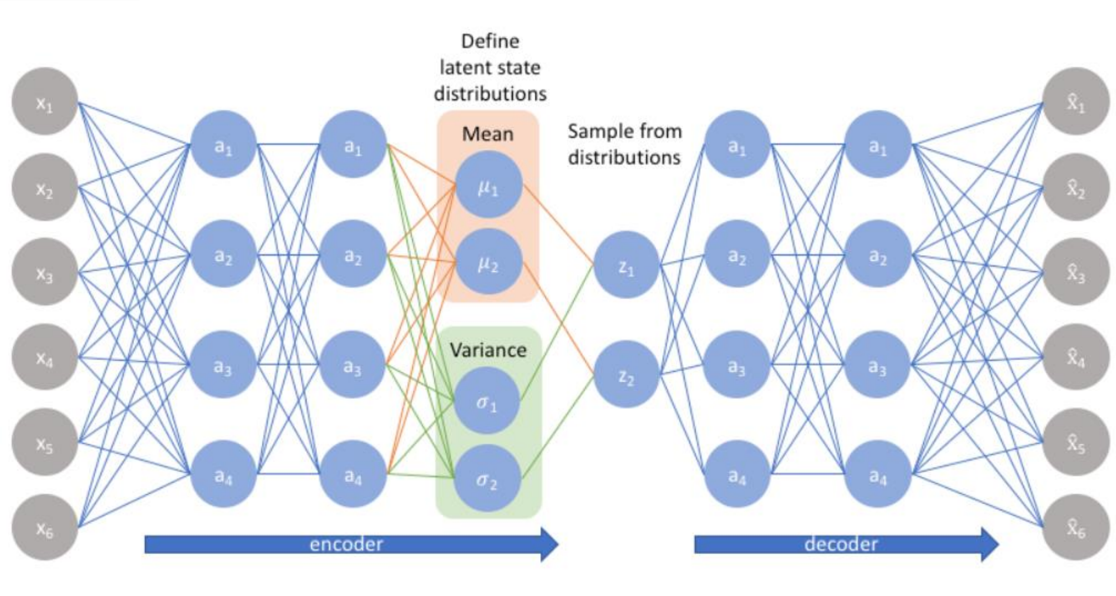

[Iten et al, PRL, 2018]

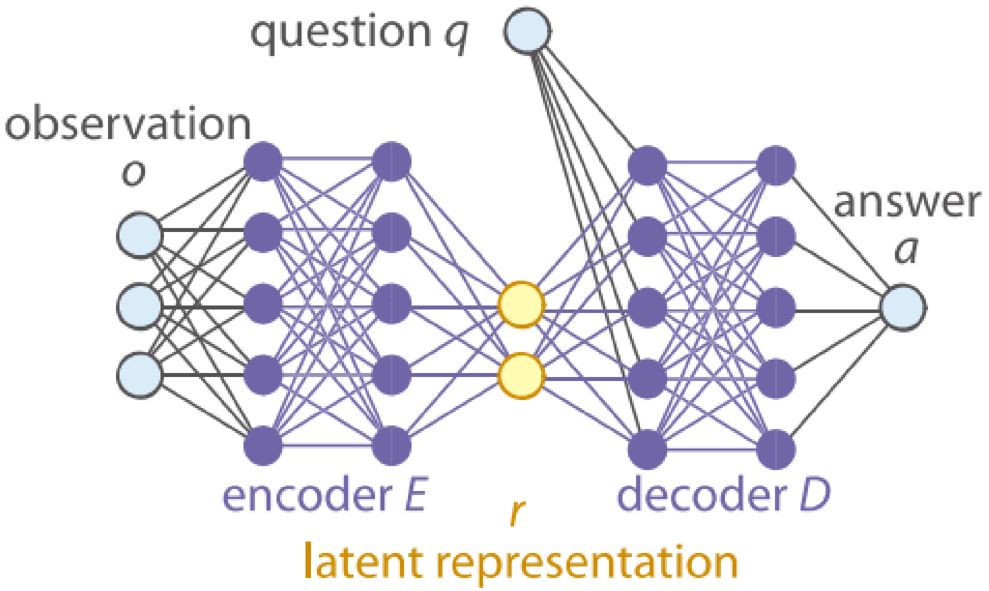

VAE: Tries to reproduce data distribution

SciNet: Tries to anwser specific questions

End of Introduction

- Questions?

Examples!

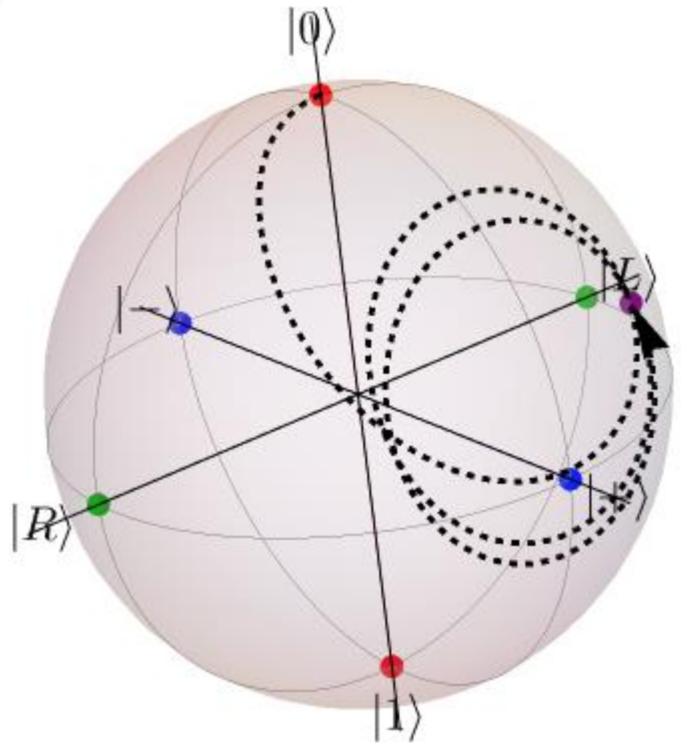

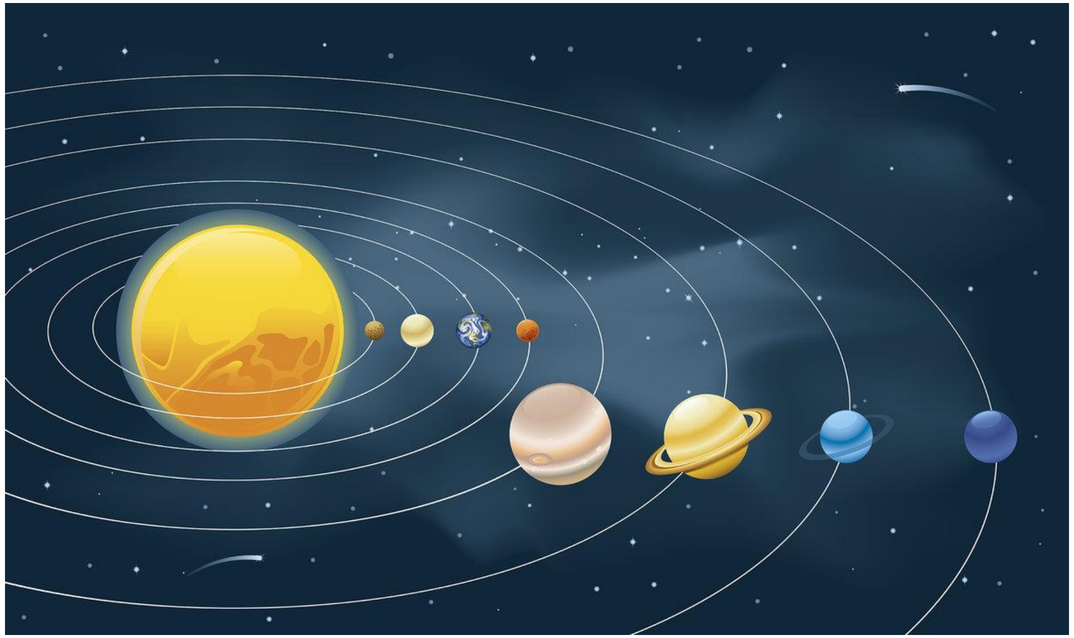

Example: Quantum State Tomography

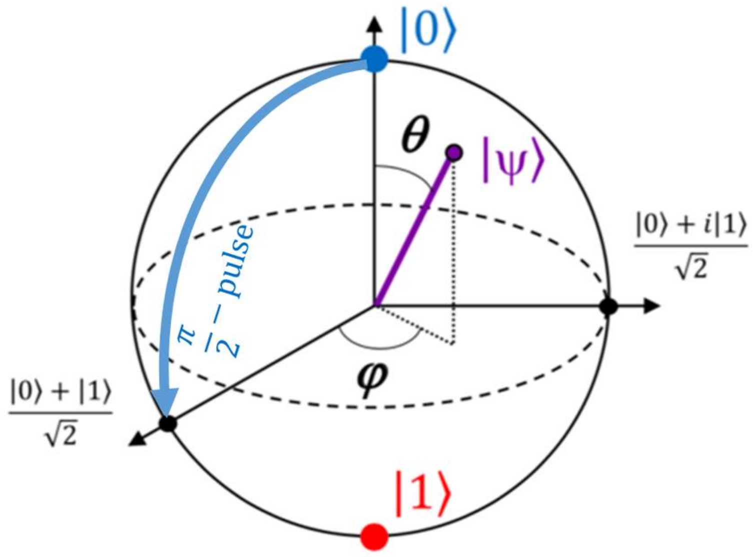

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\varphi}\sin\frac{\theta}{2}|1\rangle$$

Q: How many real parameters are needed to characterize the state?

Example: Quantum State Tomography

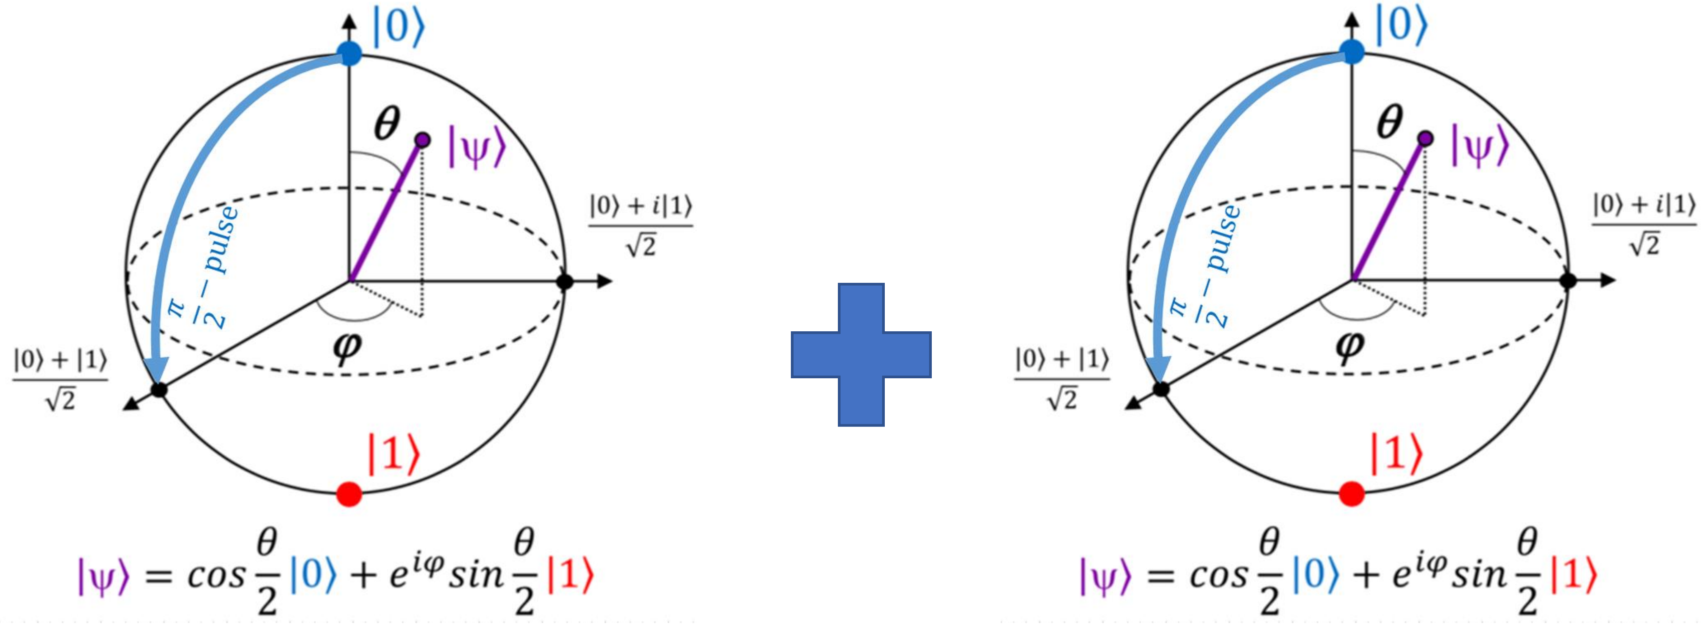

Q: How many real parameters are needed to characterize the state?

Example: Quantum State Tomography

A: 1-qubit: 2 real parameters; 2-qubit: 6 real parameters

Note: in general: $2 \cdot 2^n - 2$

Complex coeff. qu“bit“ normalization + global phase

Example: Quantum State Tomography

Task Description:

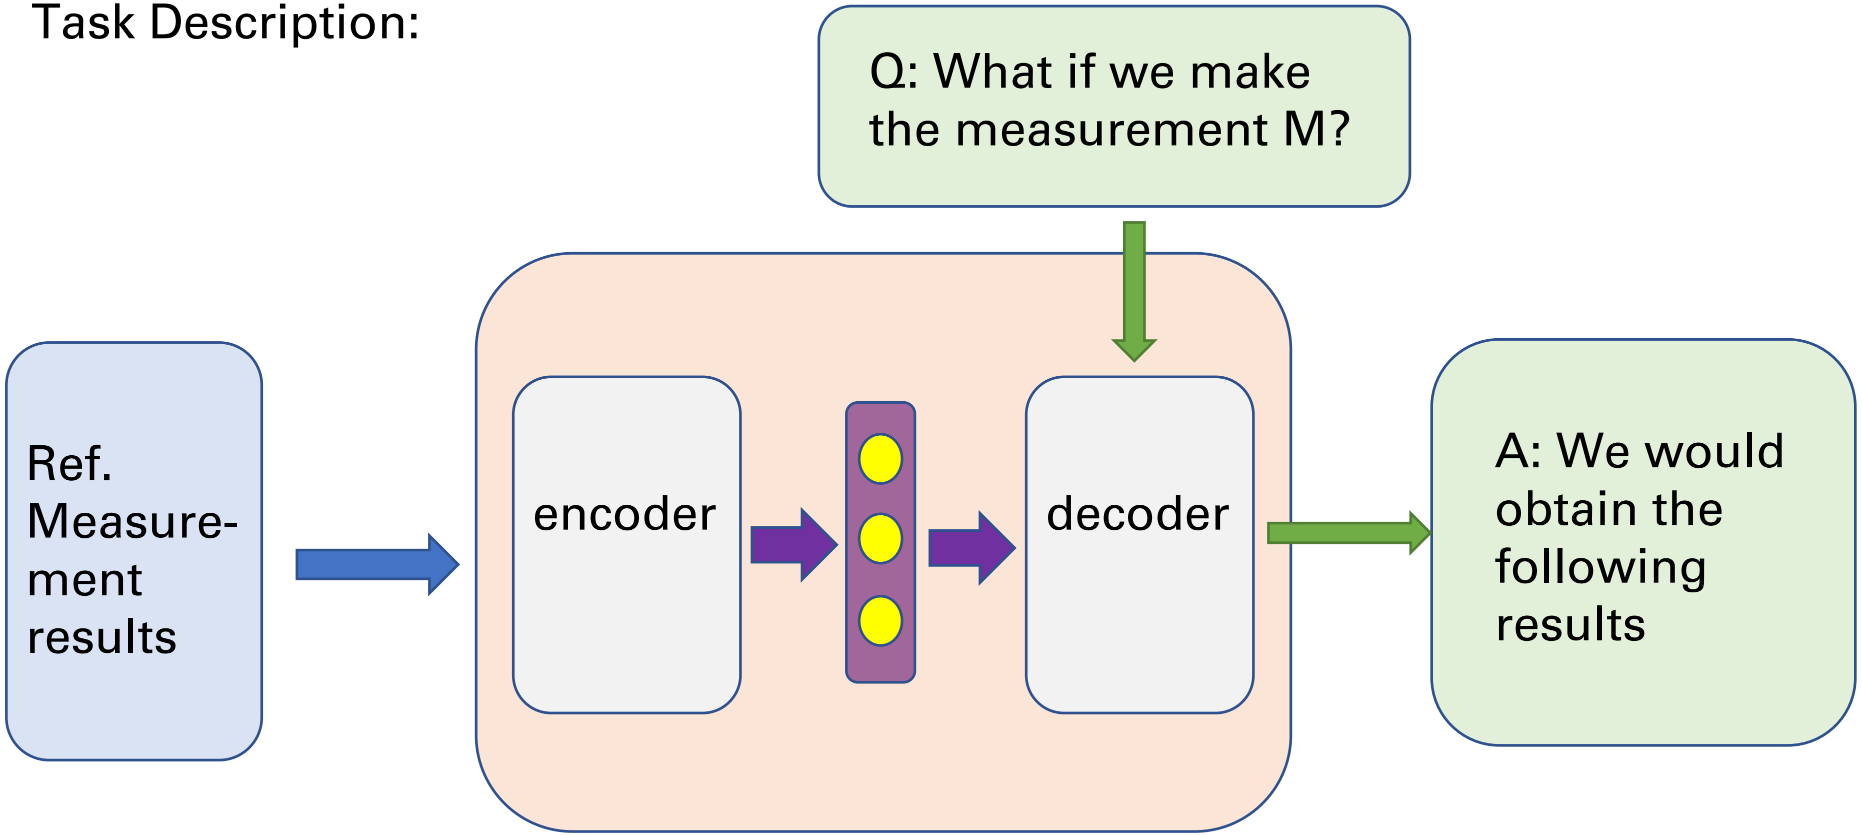

Example: Quantum State Tomography

Machine Performance:

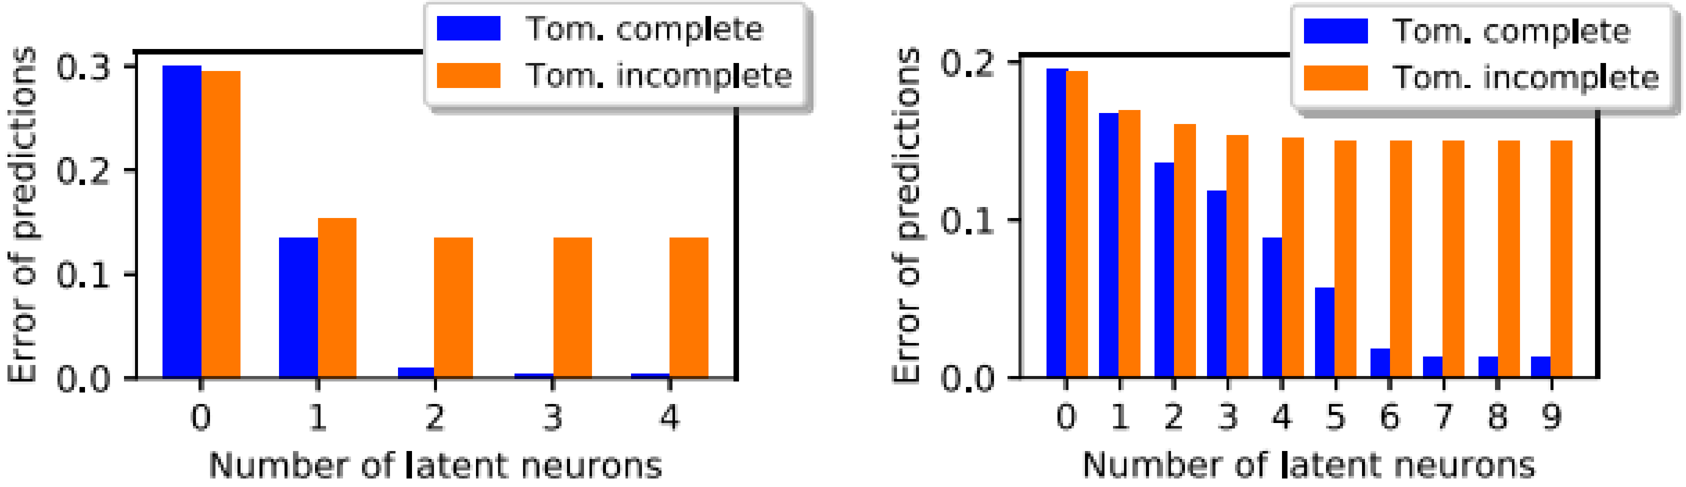

(a) One qubit (b) Two qubits

[Iten et al, PRL, 2018]

Example: Quantum State Tomography

Machine Performance:

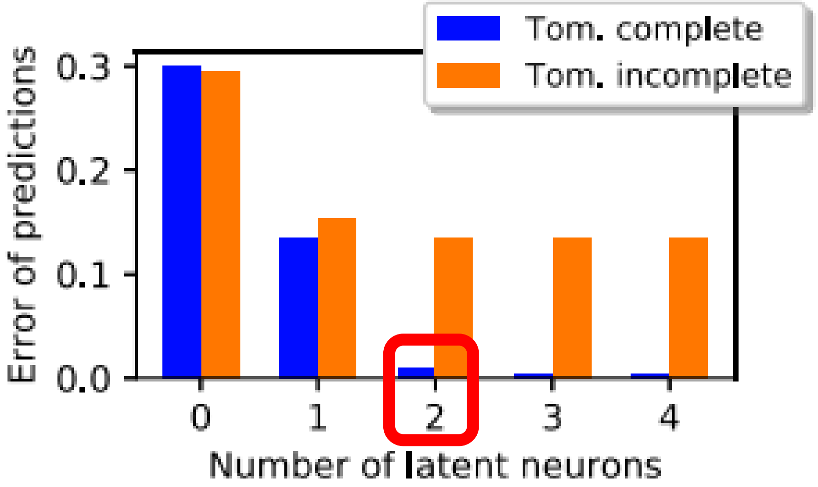

(a) One qubit

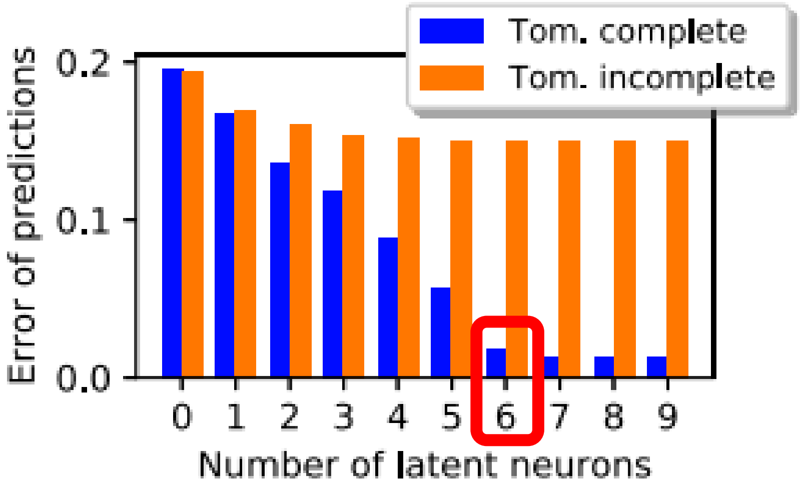

(b) Two qubits

[Iten et al, PRL, 2018]

Time-Varying Parameters

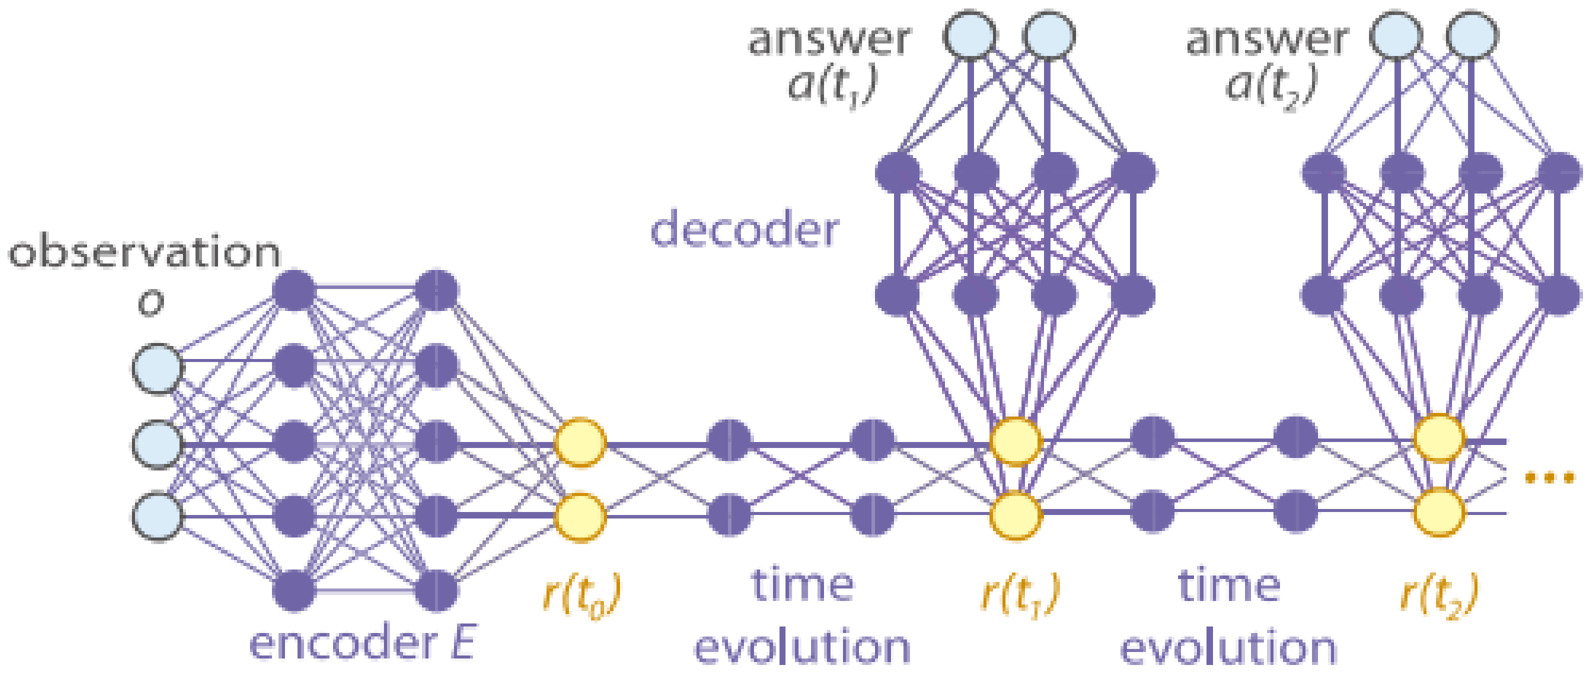

[Iten et al, PRL, 2018]

Time-Varying Parameters

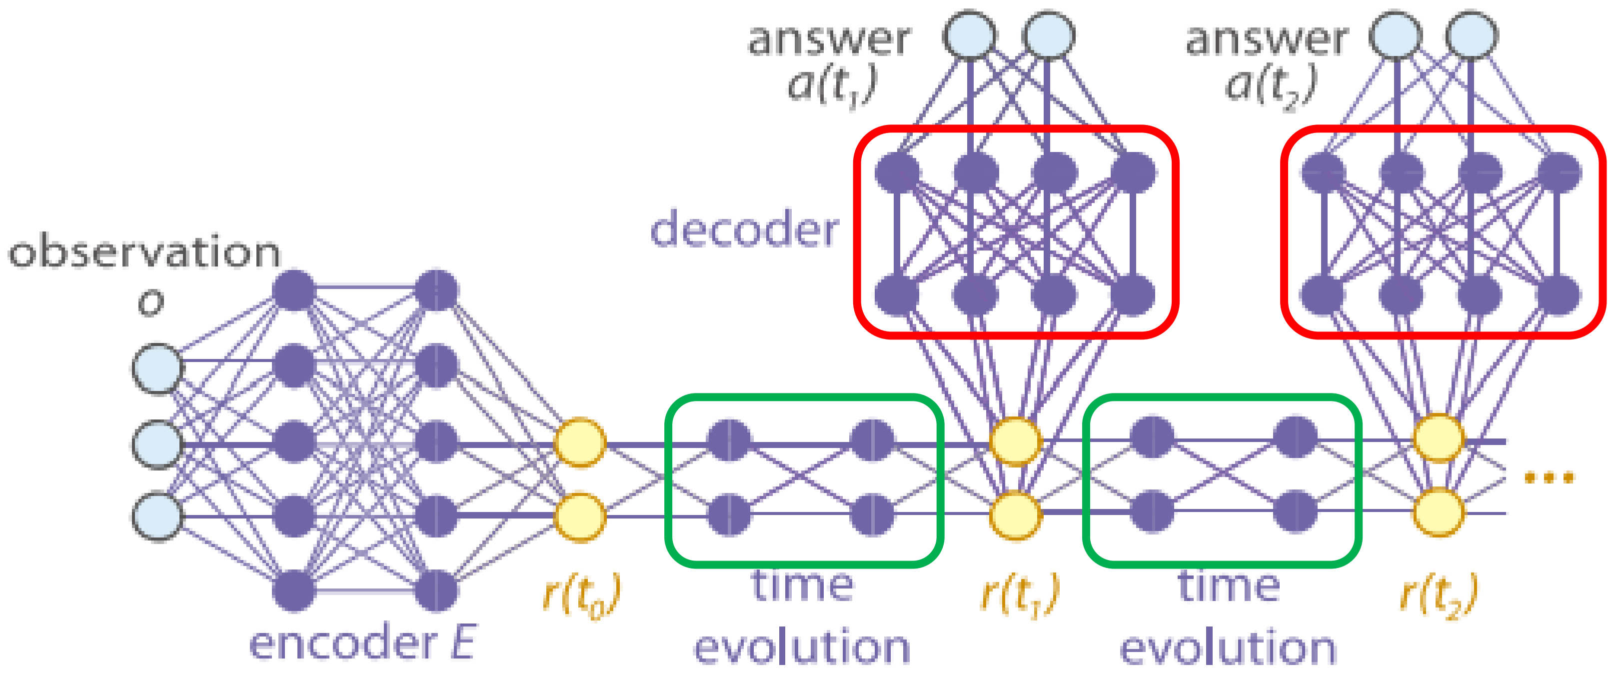

[Iten et al, PRL, 2018]

Example: Heliocentric Model

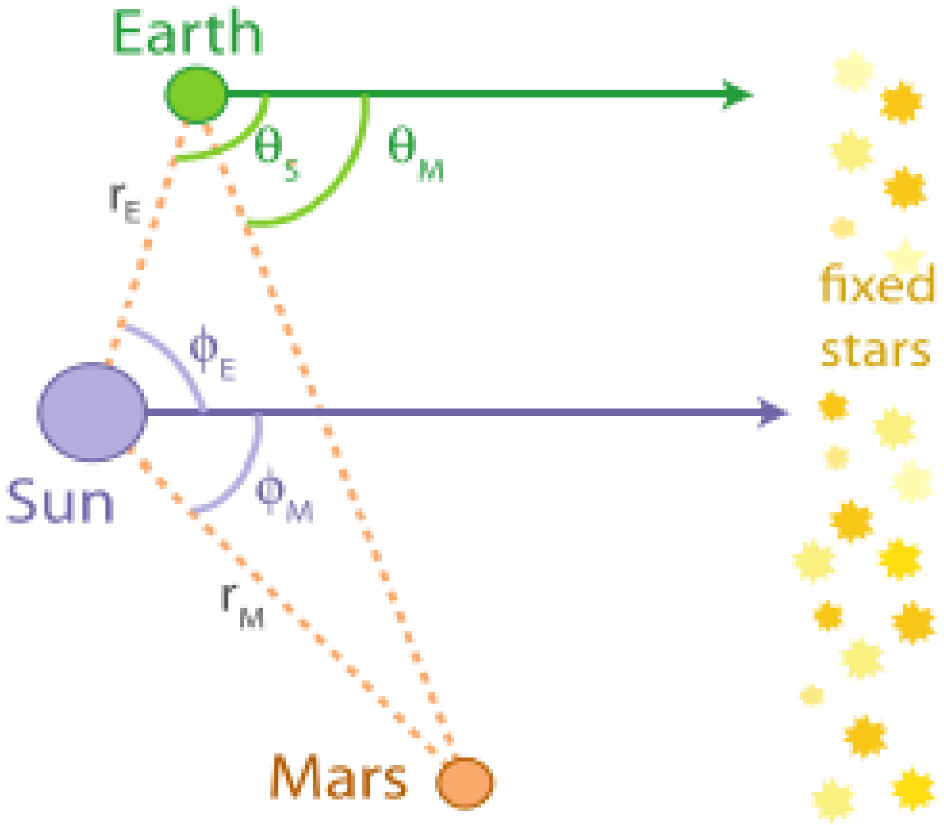

[Iten et al, PRL, 2018]

Example: Heliocentric Model

Good Parameters! (simple time evolution)

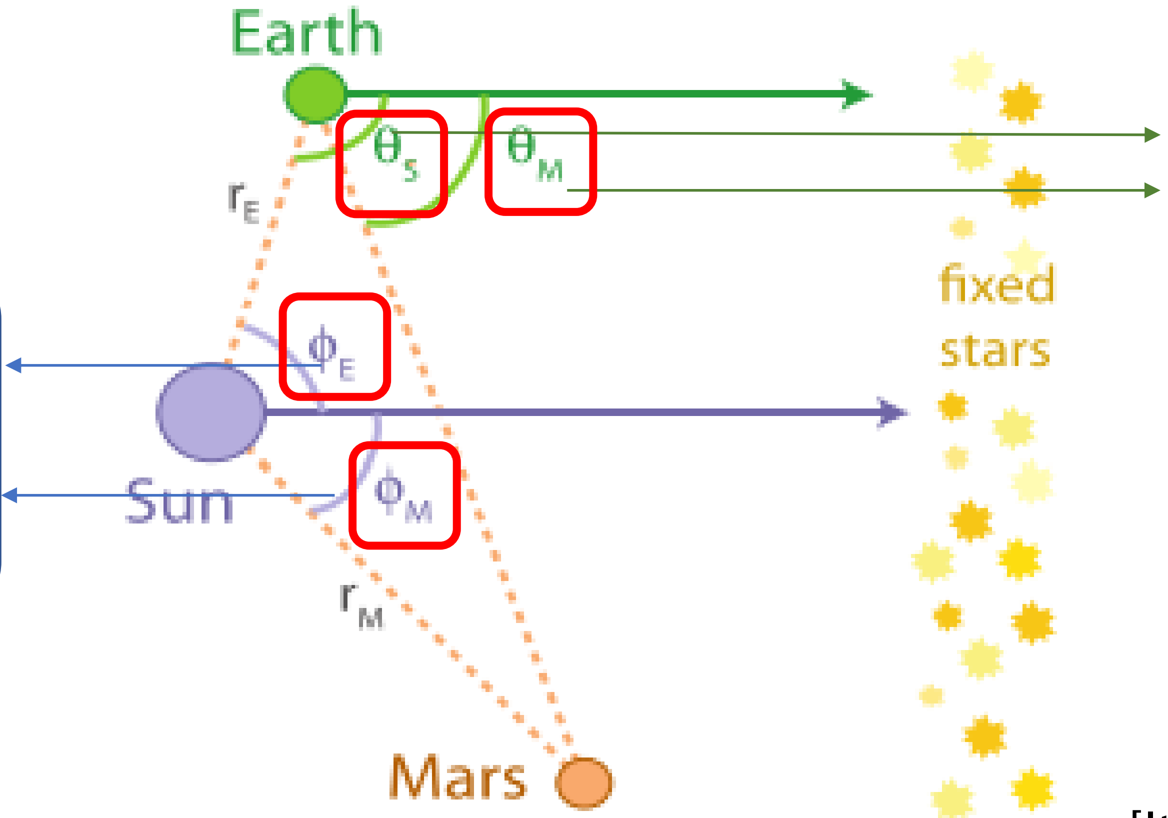

Bad Parameters! (Complicated time evolution)

[Iten et al, PRL, 2018]

Example: Heliocentric Model

[Iten et al, PRL, 2018]

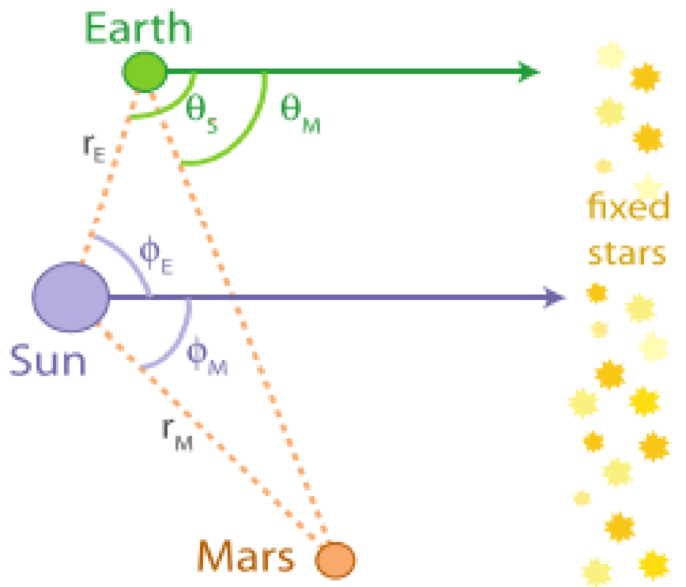

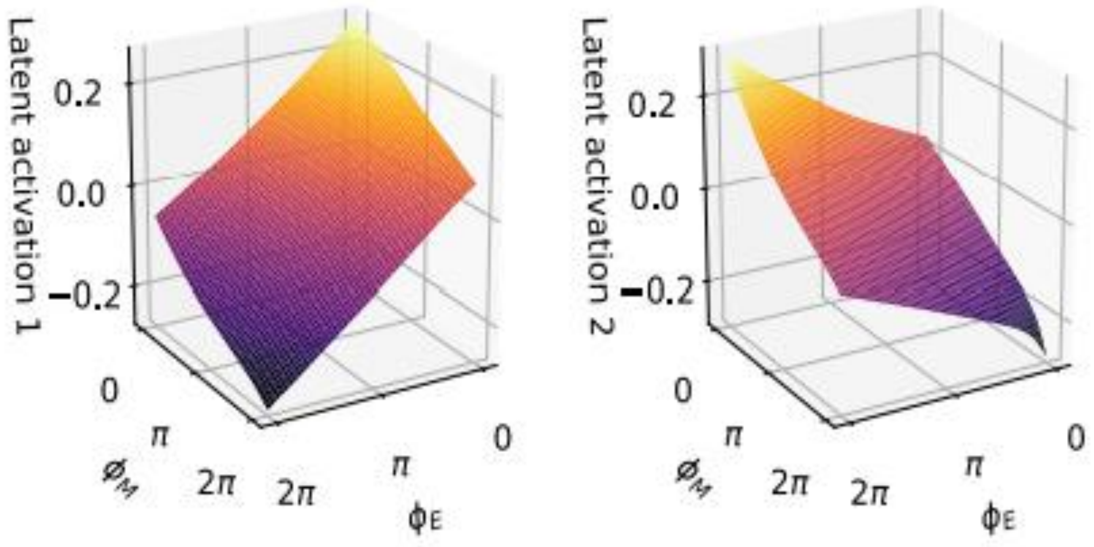

predicts the angles of Mars and the Sun with a root mean square error below 0.4% (with respect to 2π)

Example: Heliocentric Model

[Iten et al, PRL, 2018]

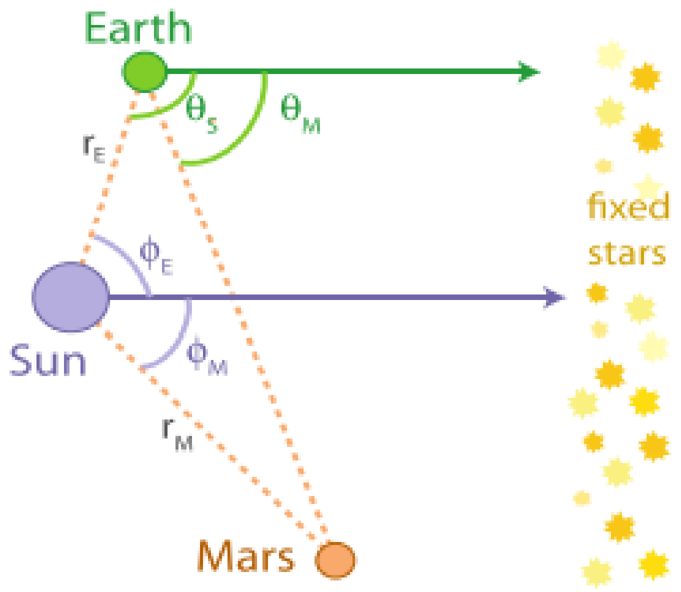

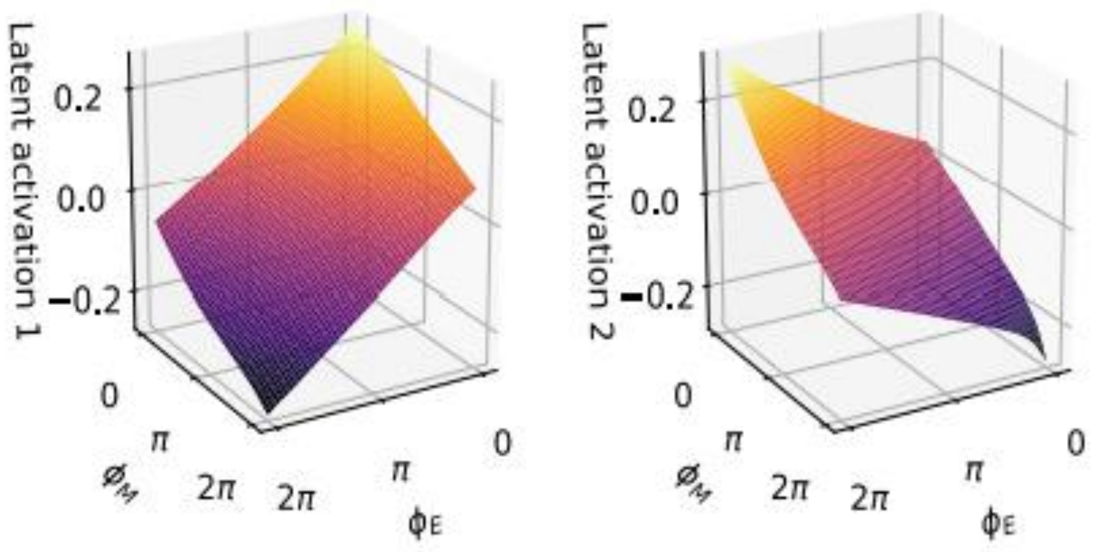

predicts the angles of Mars and the Sun with a root mean square error below 0.4% (with respect to 2π)

Activation of neurons at t_0; stores linear combinations of ϕ_M and ϕ_E

Remark

Minimal number of latent neurons ⟷ Relevant # DOF required to answer all questions

Remark

Minimal number of latent neurons

↔

Relevant # DOF required to answer all questions

Can be proven using methods from differential geometry!

End of Examples

- Questions?

Mathematics!

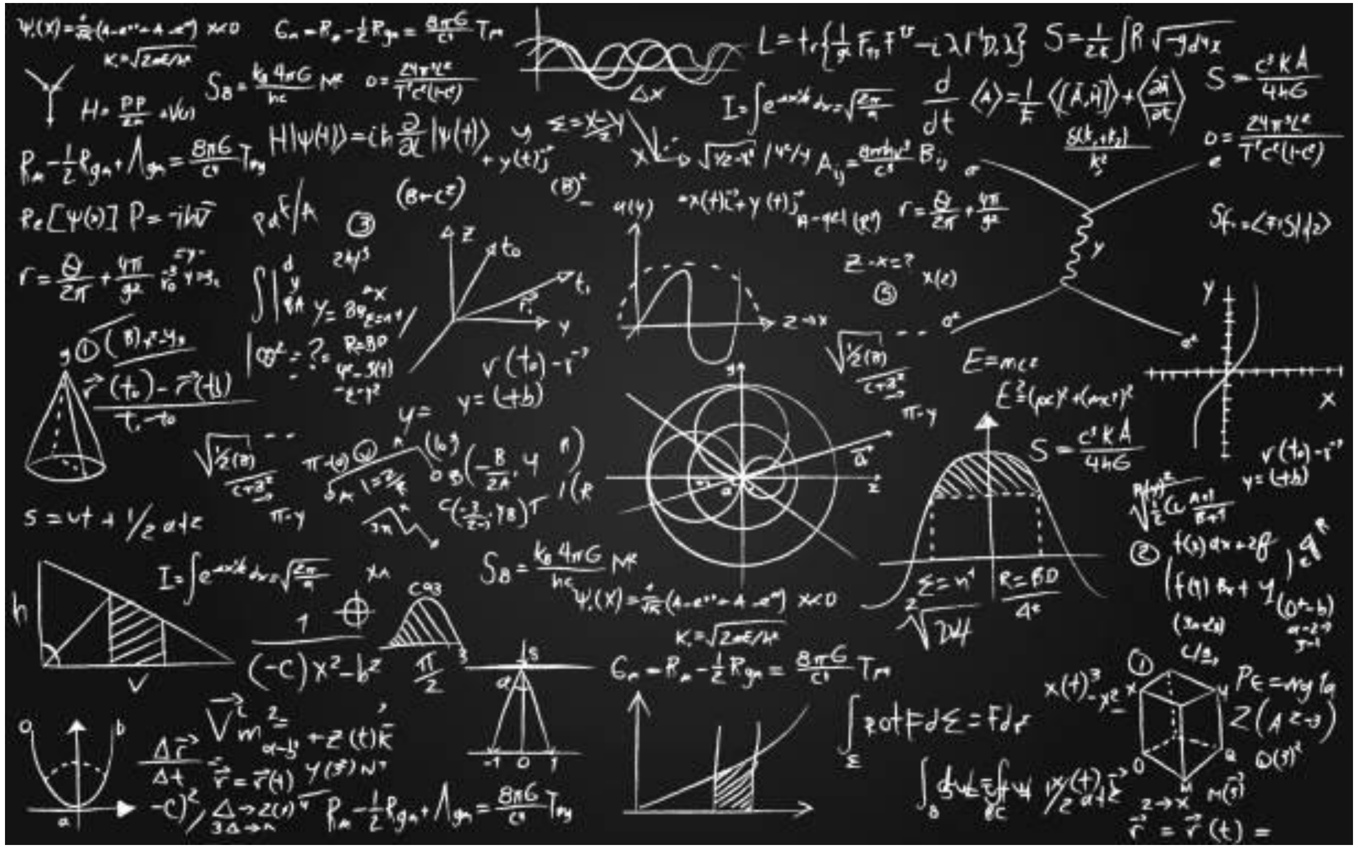

Mathematics!

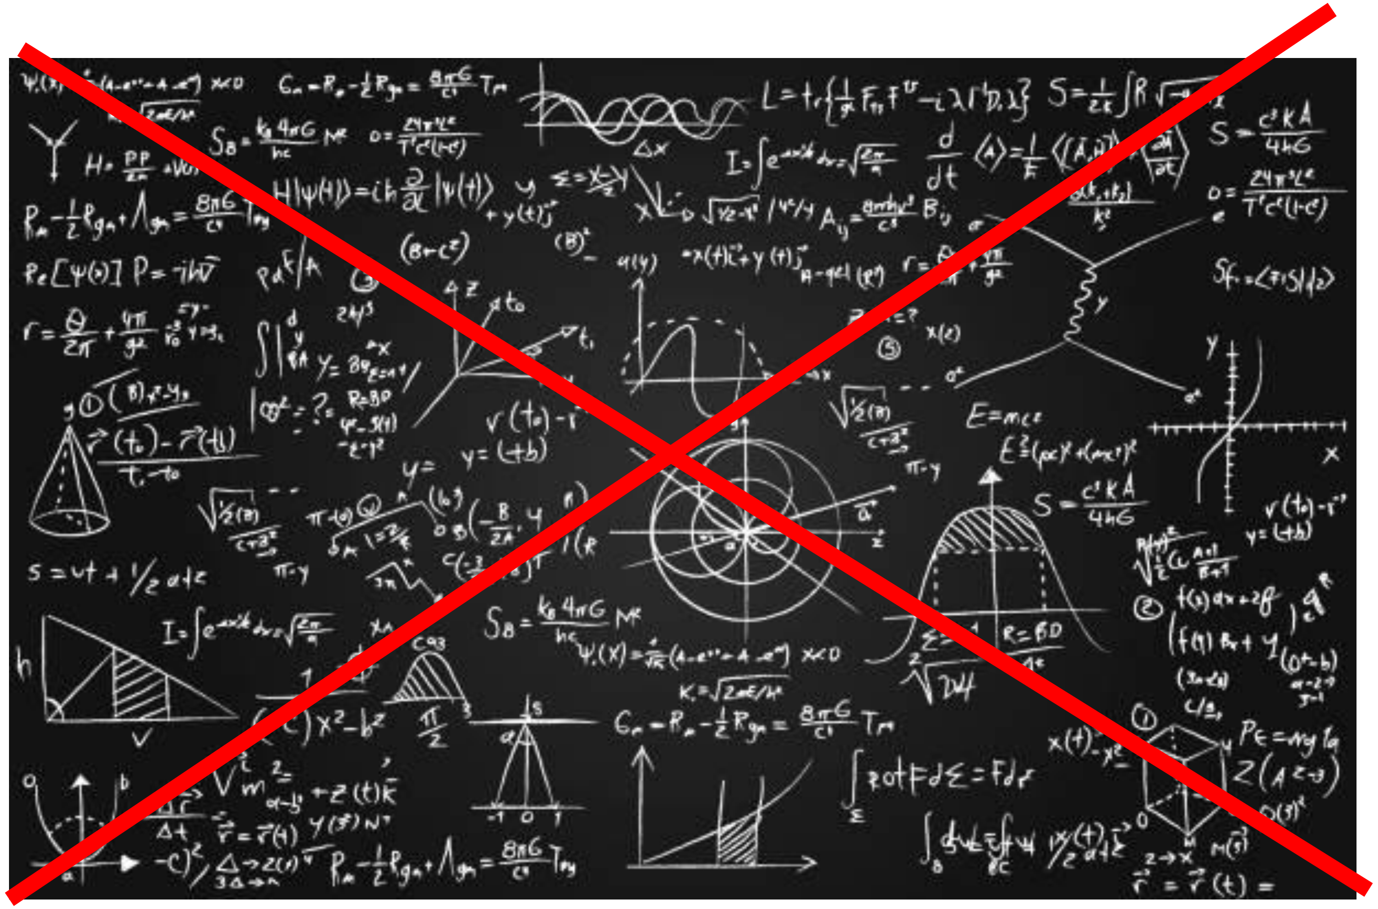

Mathematics!

$$p(z|x) = \frac{p(x|z) \cdot p(z)}{p(x)}$$

Bayesian Inference

$$\mathcal{L}_{\theta,\phi}(x) = \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\right)\right]$$

Evidence Lower Bound (ELBO)

Bayesian Inference

likelihood

Prior distribution

True posterior distribution

$$p(z|x) = \frac{p(x|z) \cdot p(z)}{p(x)}$$

Marginal likelihood

$$p(x) = \int p(x, z) dz$$

Bayesian Inference

Prior distribution

likelihood

True posterior distribution

$$p(z|x) = \frac{p(x|z) \cdot p(z)}{p(x)}$$

Marginal likelihood

$$p(x) = \int p(x,z)dz$$

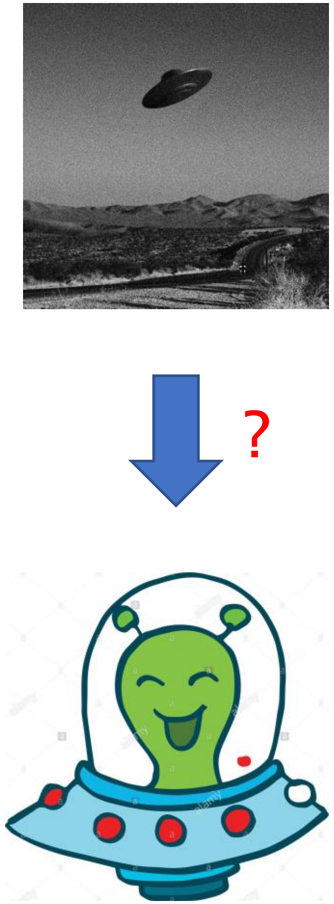

Graphical Representation

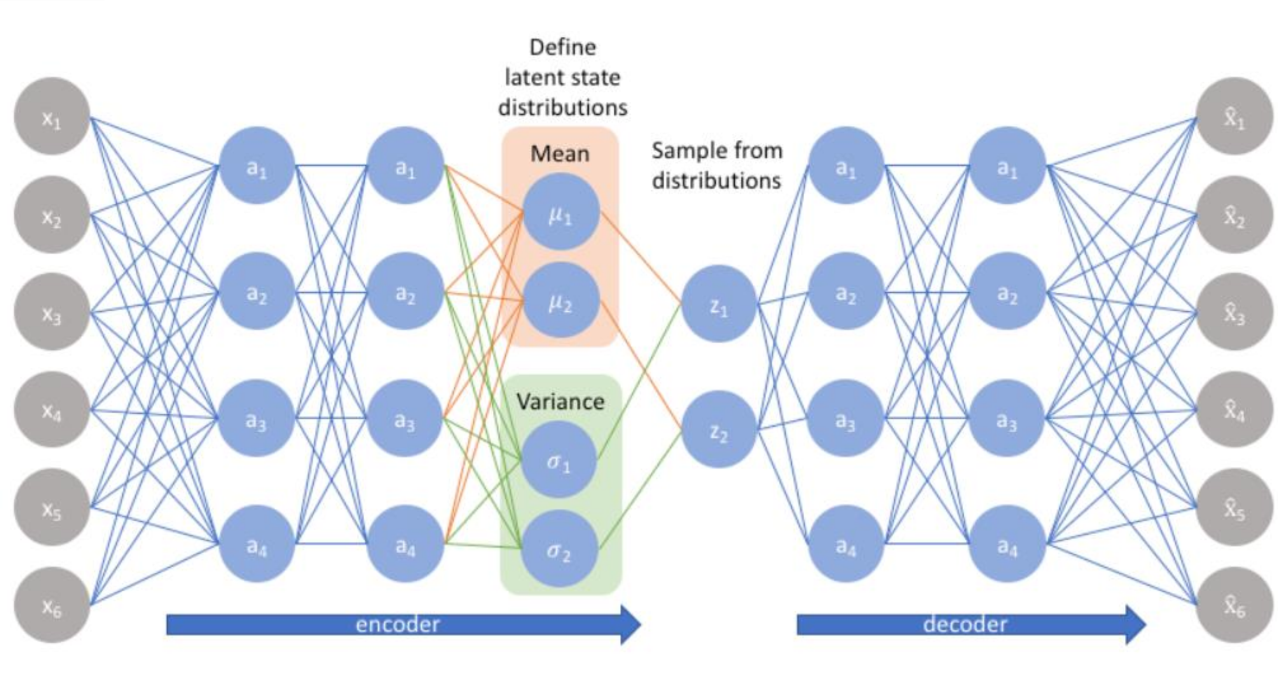

Independent Gaussians

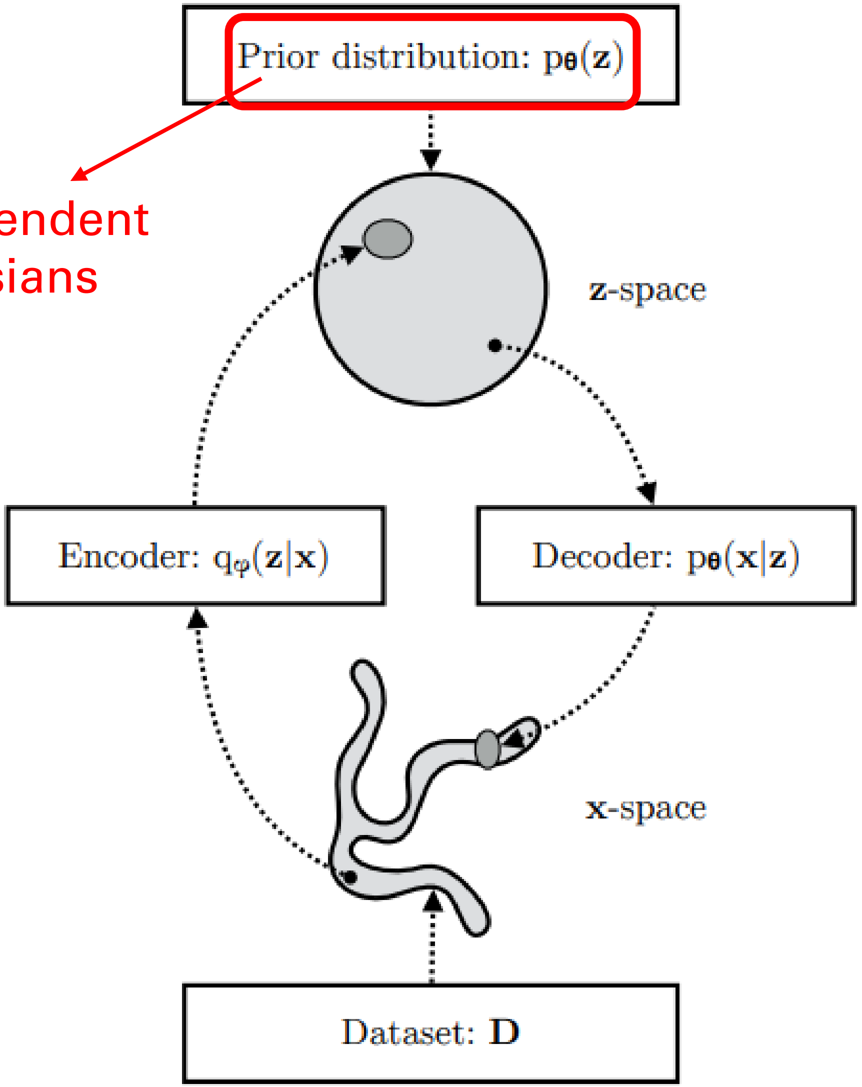

Arxiv:1906.02691v3

Graphical Representation

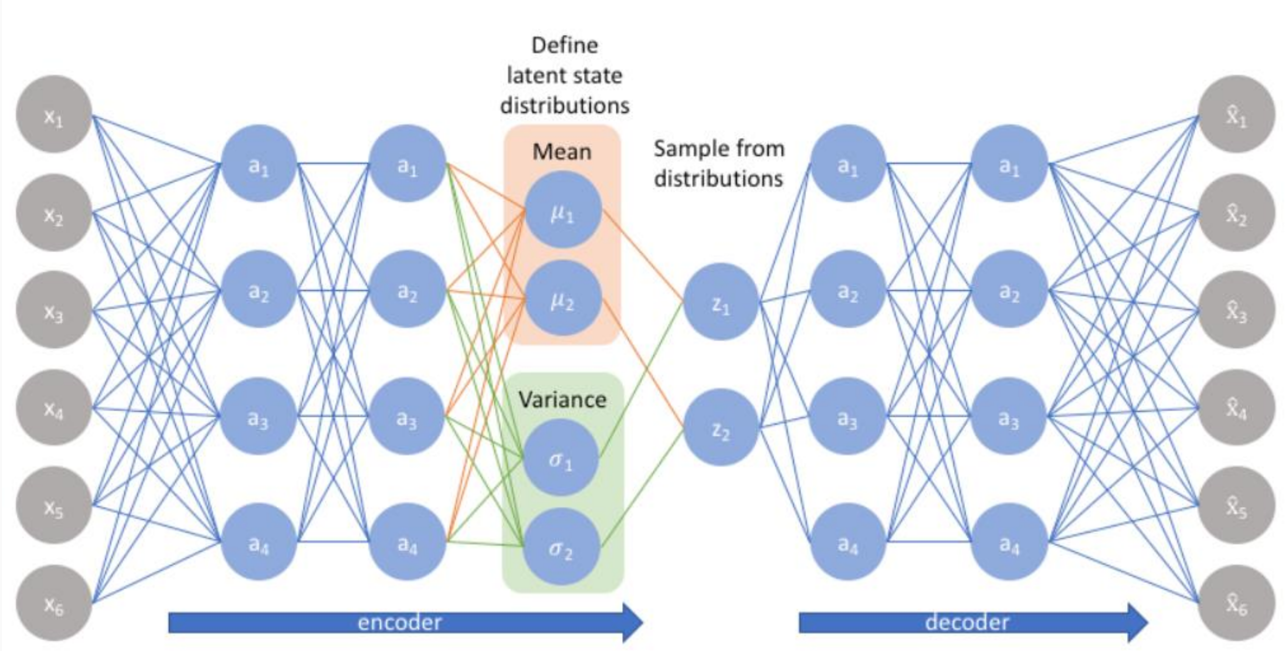

Independent Gaussians

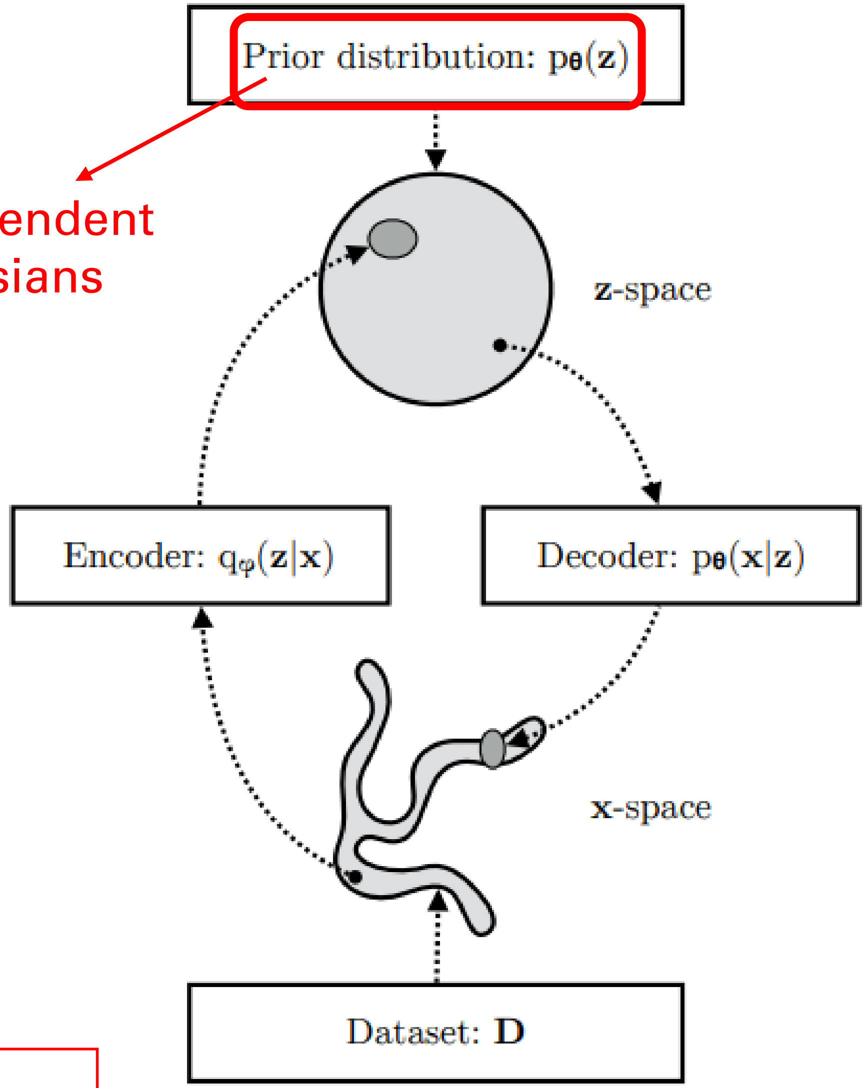

Remark: data set = weird distribution,
Training autoencoder = find variables z with simple distribution
z = equivalent to the dataset, capturing essential features

Arxiv:1906.02691v3

Variational Problem

Let‘s make the inferred model close to the true model!

Hmm...vary the inferred model to minimize some „distance measure“

Choose the KL divergence to be our „distance measure“ ☺

But we don‘t know the true model (posterior distribution)! ☹☹☹

Variational Problem

Hmm…vary ~~the inferred model~~ $\boldsymbol{q_\phi(z|x)}$ to minimize some ~~„distance measure“~~ functional $\boldsymbol{S[q_\phi(z|x)]}$

Choose the KL divergence as our ~~„distance measure“~~ functional ☺

$$D_{KL}\big(q_\phi(z|x)\big|\big|p_\theta(z|x)\big)$$

But we don‘t know the ~~true model~~) $\boldsymbol{p_\theta(z|x)}$! ☹☹☹

Let‘s make ~~the inferred model~~ $\boldsymbol{q_\phi(z|x)}$ close to ~~the true model~~ $\boldsymbol{p_\theta(z|x)}$!

Variational Problem

„The way out“

$$\mathcal{L}_{\theta,\phi}(x) \equiv \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\right)\right] \sim -D_{KL}\big(q_\phi(z|x)\,||\,p_\theta(z|x)\big)$$

Evidence Lower Bound
(ELBO)

* „The way out“ --> Derivation!

$$\log p_\theta(x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x)] = \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{p_\theta(z|x)}\right)\right]$$

$$= \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\frac{q_\phi(z|x)}{p_\theta(z|x)}\right)\right]$$

$$= \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\right)\right] + \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{q_\phi(z|x)}{p_\theta(z|x)}\right)\right]$$

$$= \mathcal{L}_{\theta,\phi}(x) + D_{KL}\left(q_\phi(z|x)\middle|\middle|p_\theta(z|x)\right)$$

$$\therefore \mathcal{L}_{\theta,\phi}(x) \sim -D_{KL}\left(q_\phi(z|x)\middle|\middle|p_\theta(z|x)\right)$$

Massaging the Equation…

$$\mathcal{L}_{\theta,\phi}(x) = \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\right)\right]$$

$$E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\big(q_\phi(z|x) || p(z)\big)$$

*Massaging the Equation…

$$\mathcal{L}_{\theta,\phi}(x) \equiv \mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x,z)}{q_\phi(z|x)}\right)\right]$$

$$\mathbb{E}_{q_\phi(z|x)}\left[\log\left(\frac{p_\theta(x|z)\cdot p(z)}{q_\phi(z|x)}\right)\right]$$

$$\mathbb{E}_{q_\phi(z|x)}\left[\log\left(p_\theta(x|z)\right) - \log\left(\frac{q_\phi(z|x)}{p(z)}\right)\right]$$

$$E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\left(q_\phi(z|x)||p(z)\right)$$

Graphical Representation

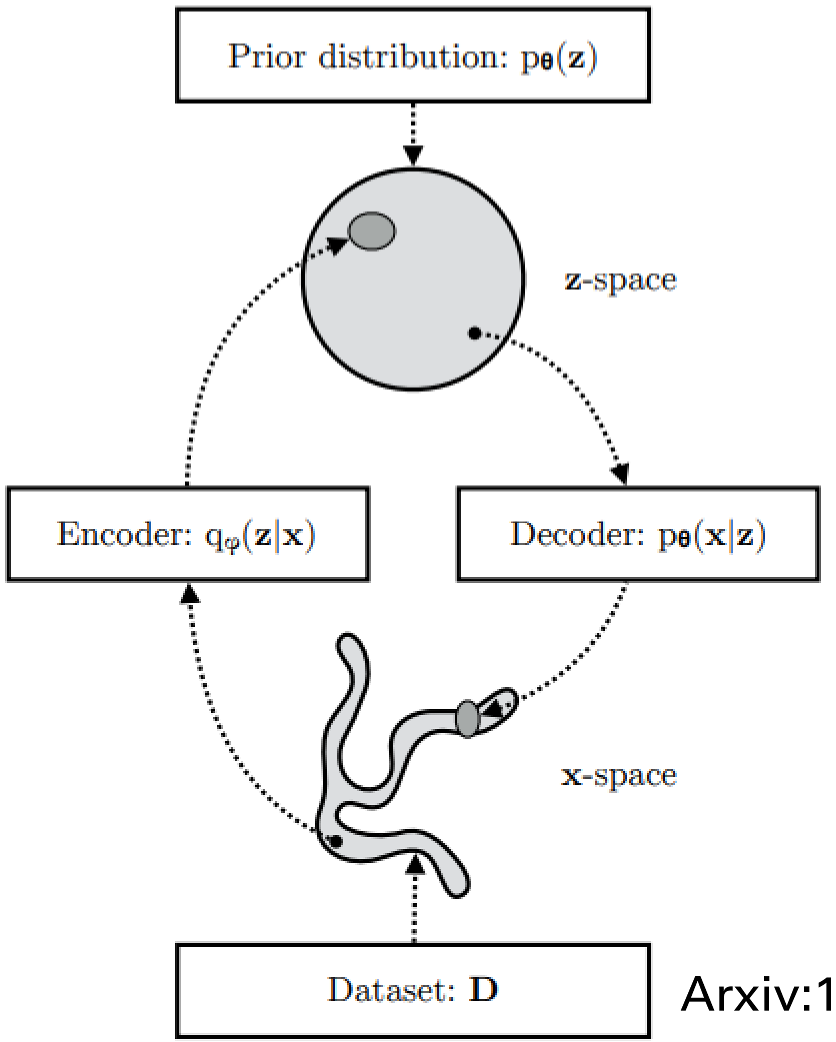

Arxiv:1906.02691v3

Optimization

- Back-propagation
- Differentiate the ELBO $\mathcal{L}_{\theta,\phi}(x)$, w.r.t. θ, ϕ
- Reminder: $\mathcal{L}_{\theta,\phi}(x) = E_{q_\phi(z|x)} \log p_\theta(x|z) - D_{KL}\left(q_\phi(z|x) || p(z)\right)$

Can't take ∇_ϕ w.r.t.
$z \sim q_\phi(z|x)$

Gradient descent
$a_{n+1} = a_n - \gamma \nabla F(a_n)$

Reparameterization Trick

- Problem…

∇_ϕ

$z \sim q_\phi(z|x)$

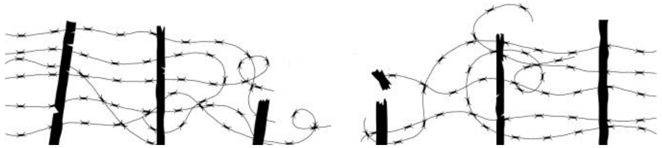

$z_i = \mu_i + \sigma_i \epsilon_i$

$\nabla_\phi \, \mathcal{L}_{\theta,\phi}(x)$

Reparameterization Trick

- Problem…

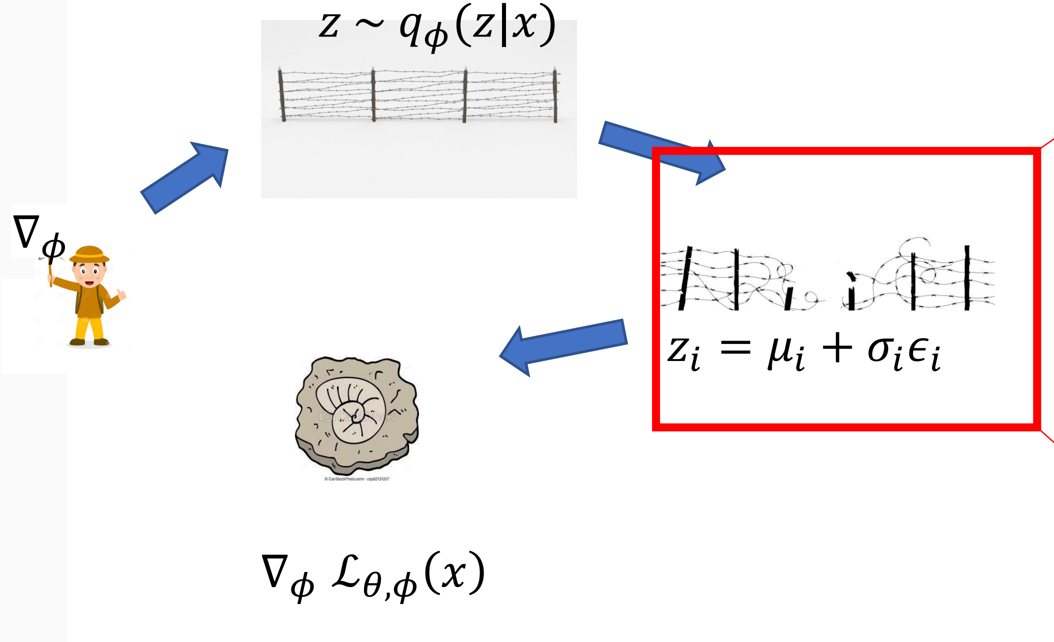

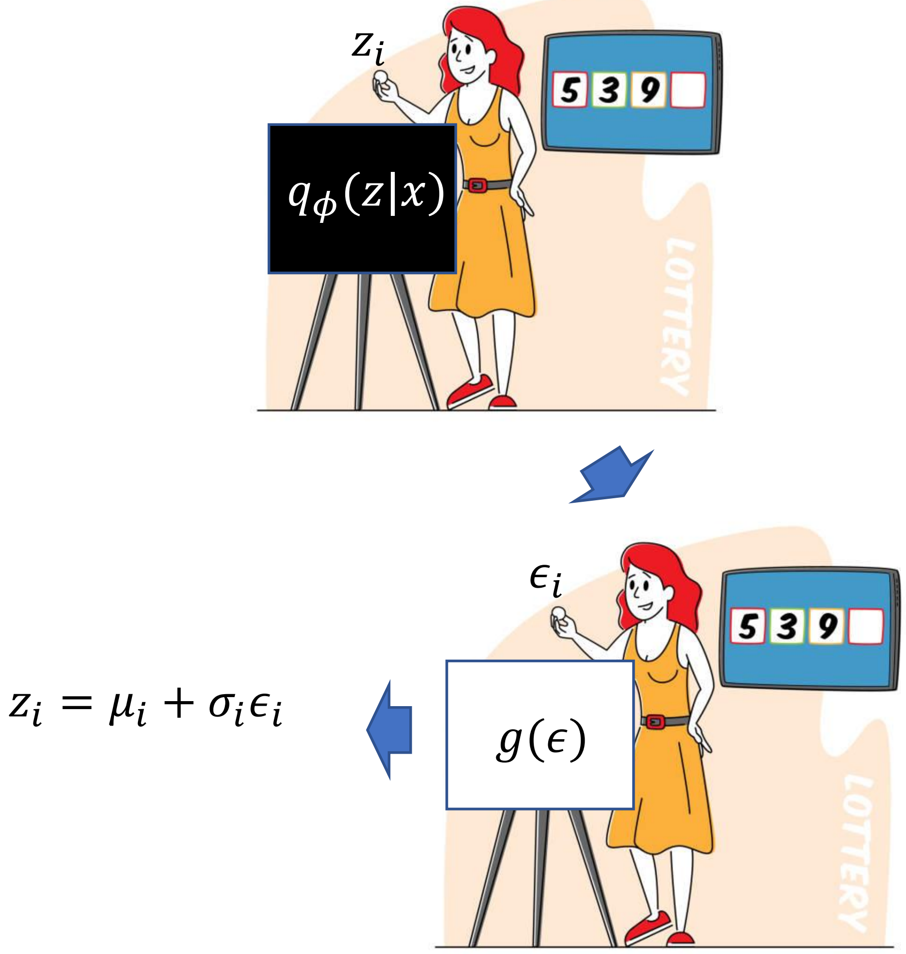

Reparameterization Trick

- Problem…

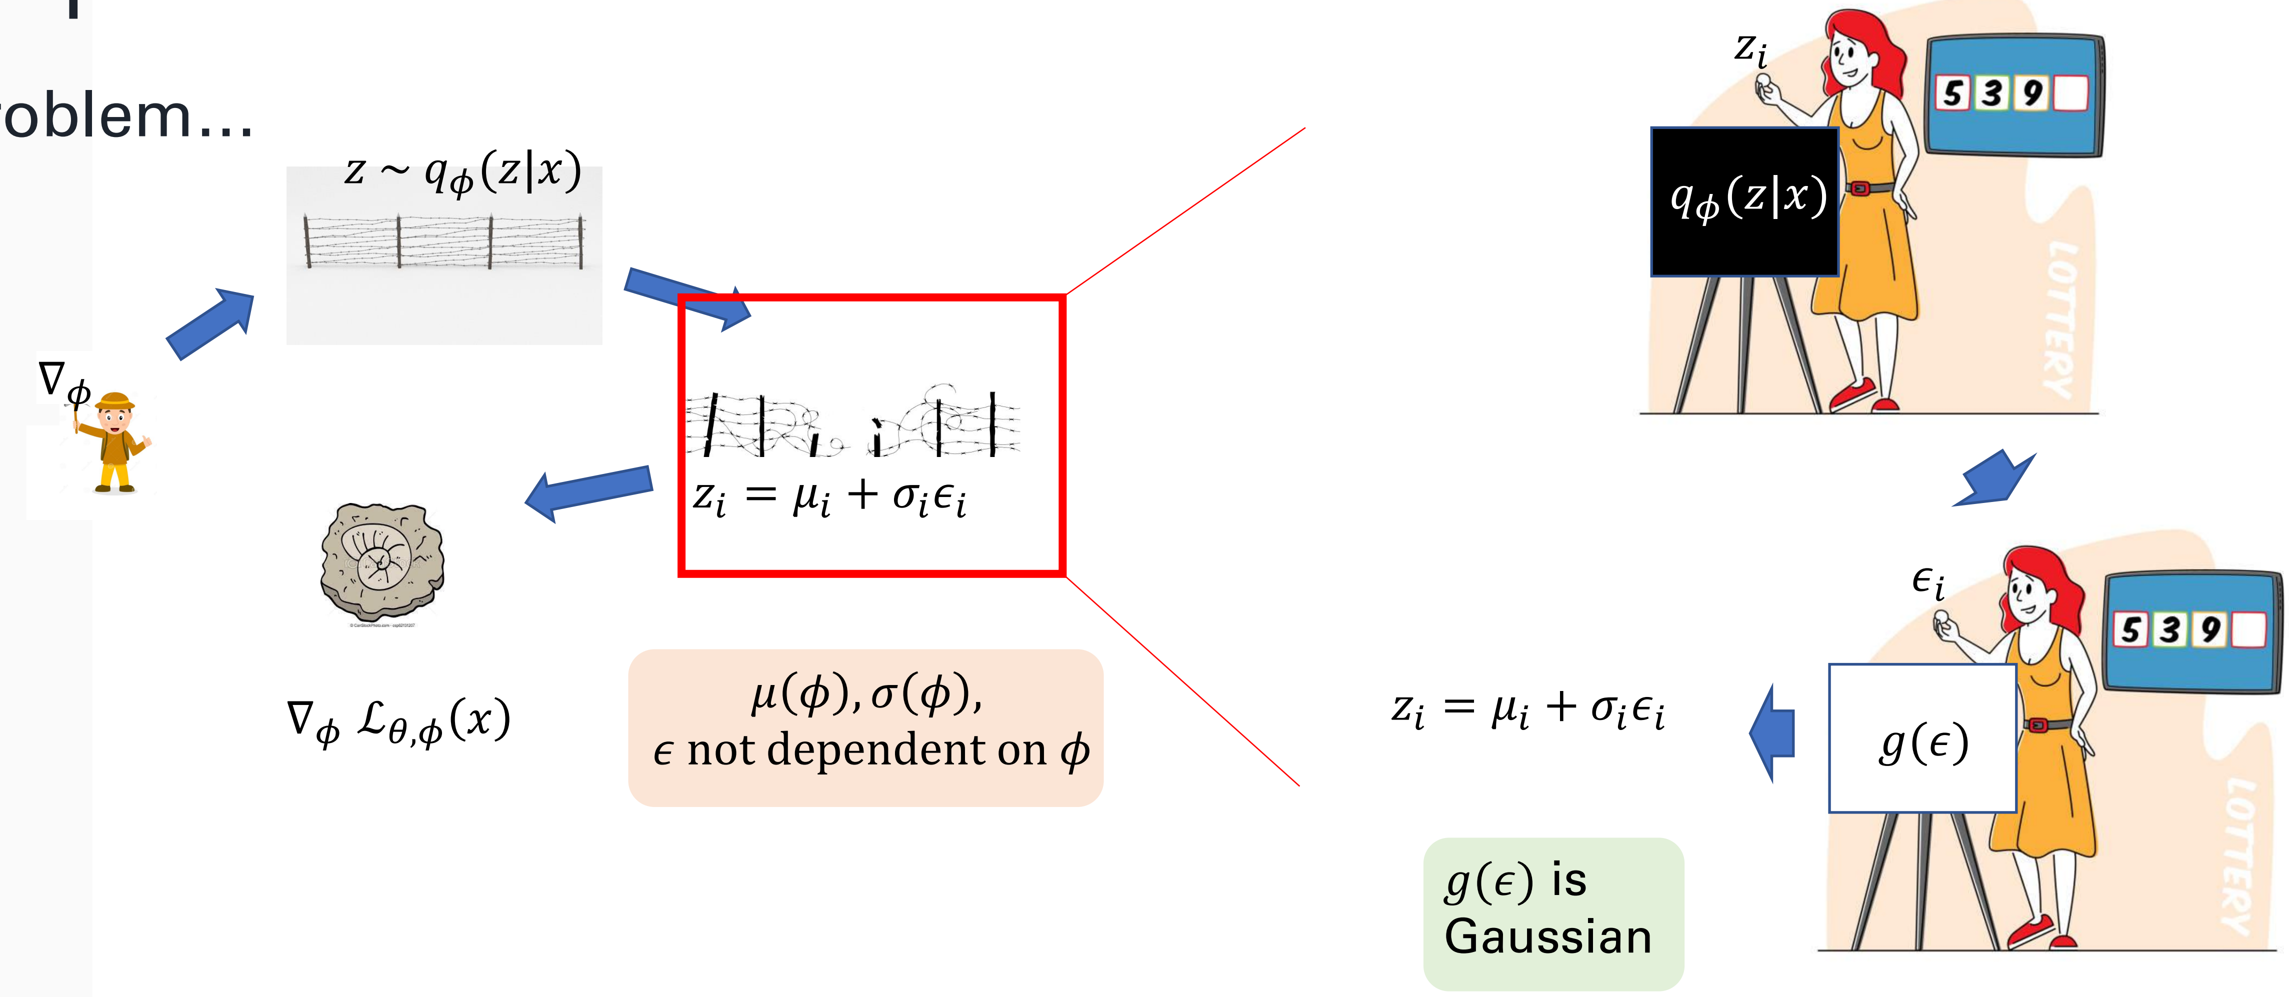

Reparameterization Trick

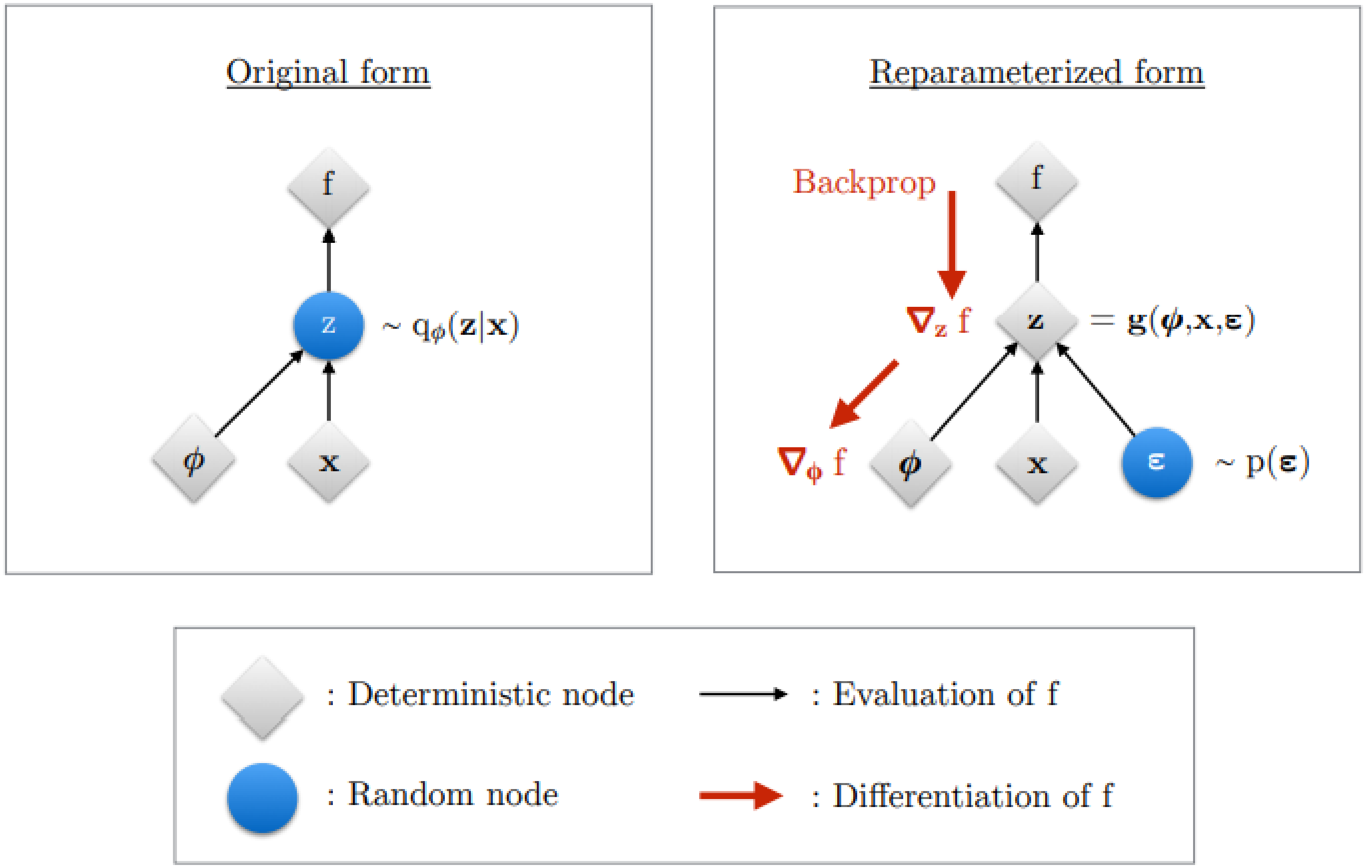

Arxiv:1906.0261v3

Recap...

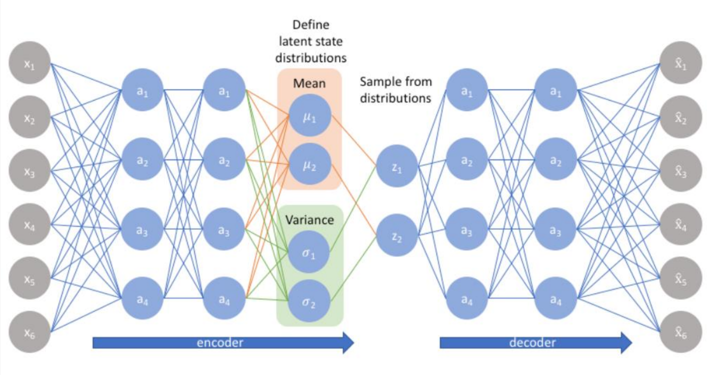

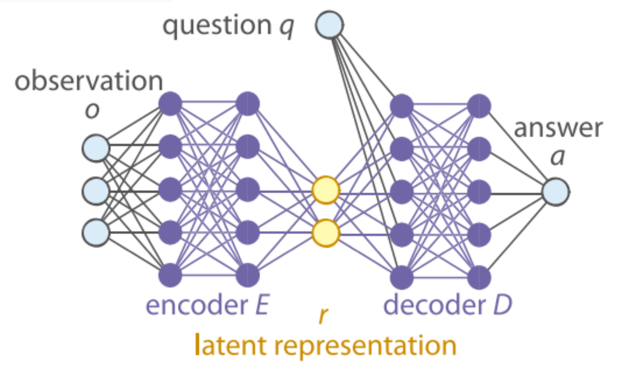

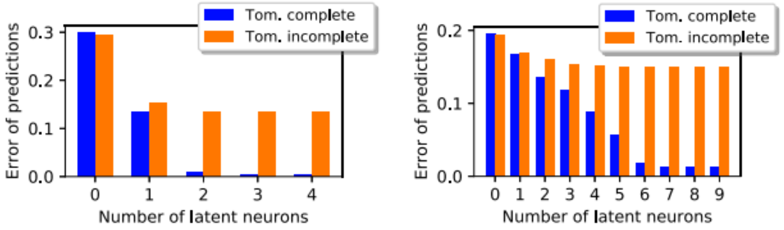

(a) One qubit (b) Two qubits

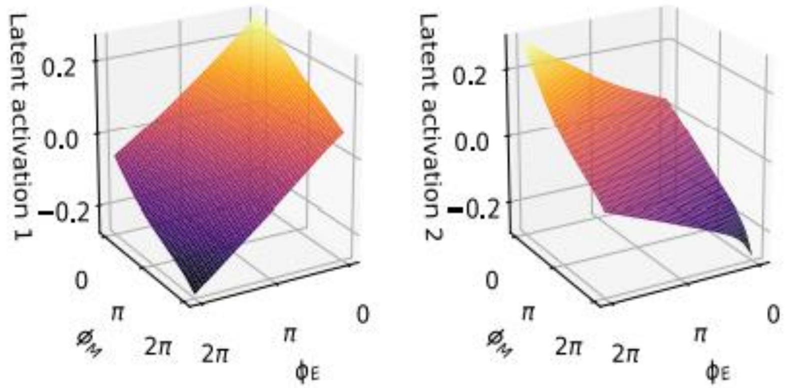

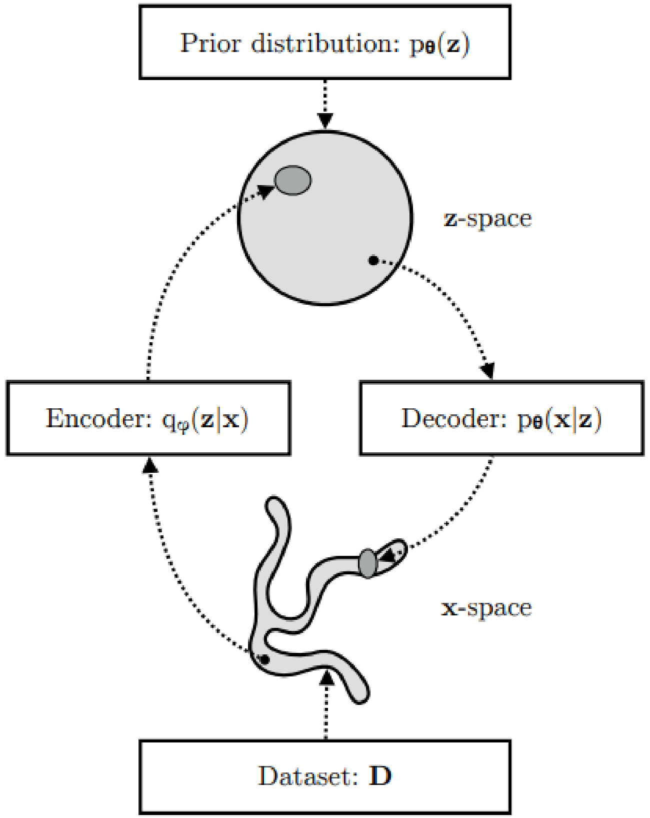

Conclusion

- VAE based SciNet structure models physical systems
- Not just the values of parameters, but also to identify relevant DOF
- Need to maximize ELBO
- Use reparameterization trick for back propagation

References

- R. Iten et al, PRL 124, 010508 (2020)
- D. Kingma and M. Welling, arxiv: 1906.02691v3, December 2019
- J. Jordan, https://www.jeremyjordan.me/variational-autoencoders/, post created March 2018, accessed July 2021