Title: Neural Networks and Quantum Mechanics

Speakers: Christian Ferko

Collection/Series: Machine Learning Initiative

Subject: Other

Date: May 23, 2025 - 2:30 PM

URL: https://pirsa.org/25050040

Abstract:

In this talk, I will survey recent developments about the connection between neural networks and models of quantum mechanics and quantum field theory. Previous work has shown that the neural network - Gaussian process correspondence can be interpreted as the statement that large-width neural networks share some properties with free, or weakly interacting, quantum field theories (QFTs). Here I will focus on 1d QFTs, or models of quantum mechanics, where one has greater theoretical control. For instance, under mild assumptions, one can prove that any model of a quantum particle admits a representation as a neural network. Cherished features of quantum mechanics, such as uncertainty relations, emerge from specific architectural choices that are made to satisfy the axioms of quantum theory. Based on 2504.05462 with Jim Halverson.

Neural Networks and Quantum Mechanics

Based on 2504.05462 with Jim Halverson

Christian Ferko

Northeastern University and IAIFI, the
Institute for Artificial Intelligence and Fundamental Interactions

Machine Learning Initiative Seminar
Perimeter Institute for Theoretical Physics
May 23, 2025

Motivation: new tools for quantum field theory.

We still lack a suitably general framework for understanding quantum field theories. When we first meet QFTs, we define them via expressions like

$$\mathcal{Z} = \int \mathcal{D}\varphi \, \exp\left(-\frac{i}{\hbar} \int d^d x \, \mathcal{L}(\varphi) \right) ,$$

which involve a local Lagrangian density.

This is appropriate for certain free or weakly interacting QFTs. However,

1. Many (perhaps most) QFTs have no Lagrangian description.
2. Some theories are described by multiple Lagrangians due to *dualities*.
3. Even when we have a Lagrangian, one should define the measure $\mathcal{D}\varphi$.

These issues motivate us to develop new techniques to understand QFTs. One toolkit which has proven useful in other domains is *neural networks*.

Neural networks are universal approximators.

Why might neural networks be helpful? Roughly, the path integral asks us to integrate over functions. Every function can be approximated by a NN.

Theorem 1 (Cybenko)

Let $f : \mathbb{R}^d \to \mathbb{R}$ be a continuous function and fix a compact set $K \subset \mathbb{R}^d$. For any $\epsilon > 0$, there exists a neural network with a single hidden layer,

$$\phi(x) = \sum_{i=1}^{N} \sum_{j=1}^{d} w_i^{(1)} \sigma \left(w_{ij}^{(0)} x_j + b_i^{(0)} \right) + b^{(1)} ,$$

where $\sigma : \mathbb{R} \to \mathbb{R}$ is a non-polynomial activation function such that

$$\sup_{x \in K} |f(x) - \phi(x)| < \epsilon .$$

Can we replace a path integral with an "integral over neural networks"?

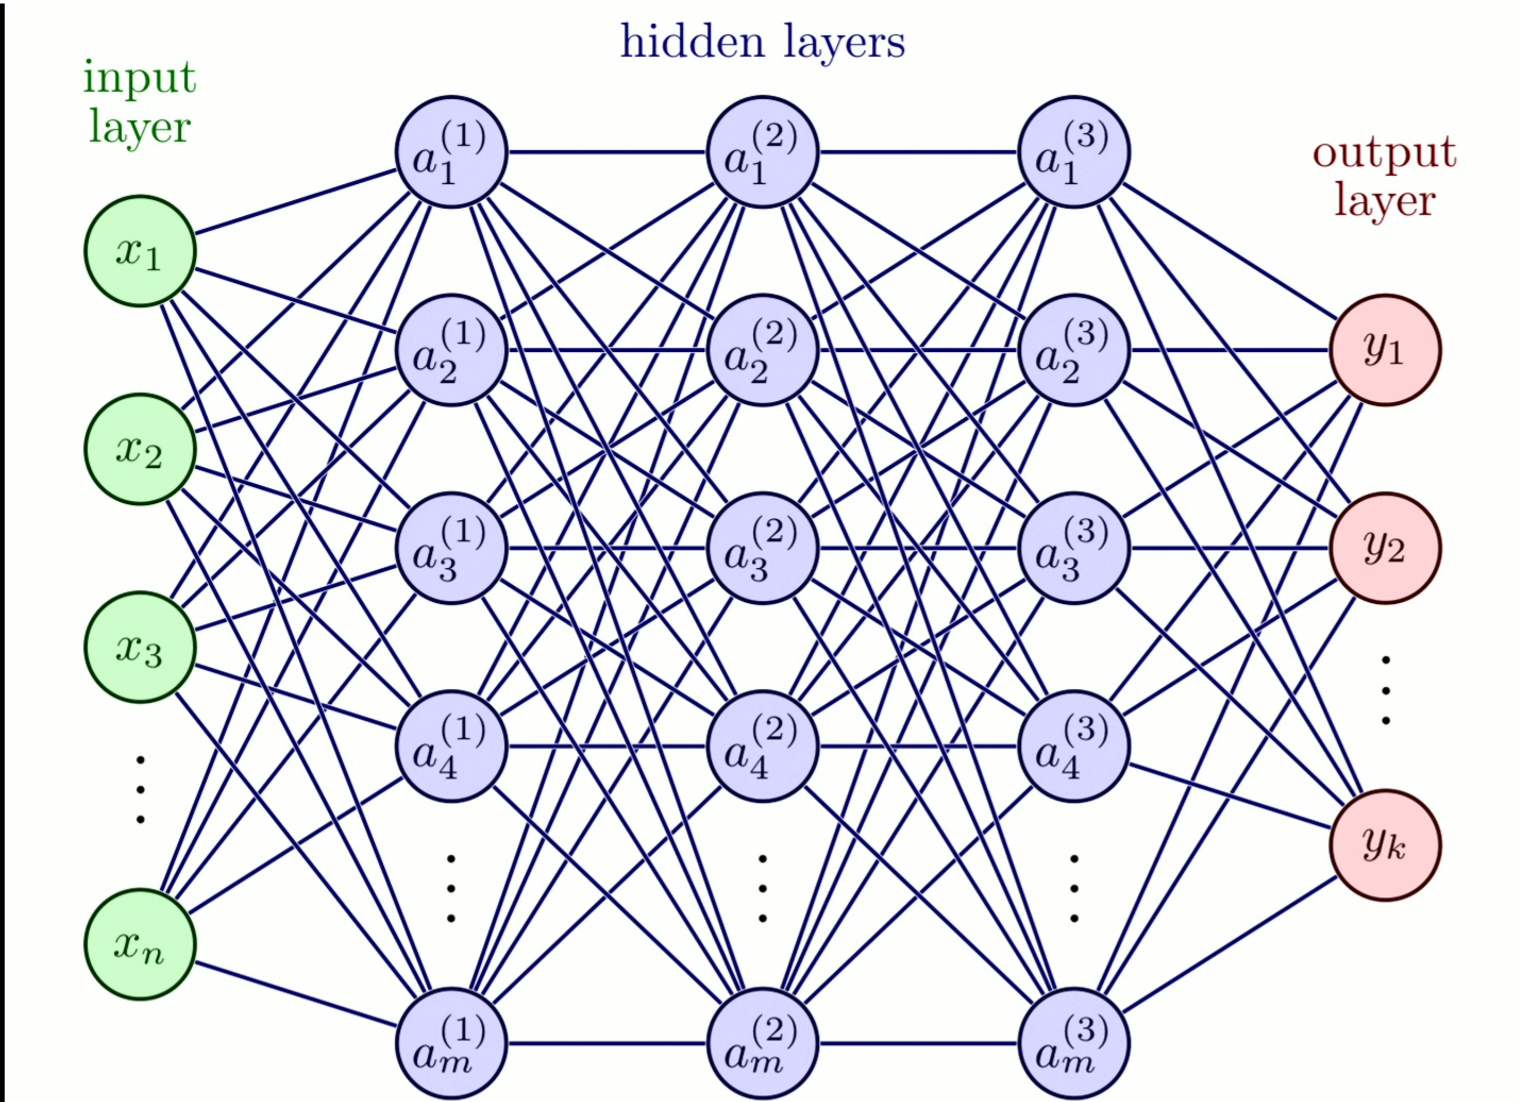
hidden layers
input layer
output layer
x_1
x_2
x_3
x_n
$a_1^{(1)}$
$a_2^{(1)}$
$a_3^{(1)}$
$a_4^{(1)}$
$a_m^{(1)}$
$a_1^{(2)}$
$a_2^{(2)}$
$a_3^{(2)}$
$a_4^{(2)}$
$a_m^{(2)}$
$a_1^{(3)}$
$a_2^{(3)}$
$a_3^{(3)}$
$a_4^{(3)}$
$a_m^{(3)}$
y_1
y_2
y_k

Three questions about neural networks.

To delineate the scope of this talk, let me point out that there are (at least) three broad sub-topics about NNs that one could investigate:

1. Expressivity: how "powerful" is a neural network, in the sense of its ability to approximate functions?
2. Statistics: what are the properties of ensembles of neural network outputs, given a particular probability distribution for the weights?
3. Training: how do neural networks evolve over training time, as their parameters are tuned to minimize some loss function?

Point (1) is the subject of Cybenko's theorem. Point (3) is interesting, but is not our focus here: we instead think of ensembles of NN with fixed parameter distributions (e.g. at initialization or after finite training time).

Point (2), and its relation to QFT, is the subject of the rest of this talk.

Roadmap.

Goal: explore the conditions under which the statistics of neural networks can reproduce the correlation functions of local operators in conventional (local, Poincaré invariant, unitary) quantum field theories.

The plan is as follows:

- ☑ Part 1: Introduction and motivation.
- ☐ Part 2: Generalities on neural network field theory.
- ☐ Part 3: New results on neural network quantum mechanics.
- ☐ Part 4: Summary and future directions.

A general notion of neural network.

For the purposes of this talk, a *neural network* is any parameterized family of functions $\phi_\theta : \mathbb{R}^n \to \mathbb{R}^m$ depending on a set of parameters θ, which are random variables drawn from a probability distribution $P(\theta)$.

A simple example of a neural network is

$$\phi_\theta(t) = \cos(t + \theta) \,, \qquad \theta \sim U[0, 2\pi] \,,$$

which means θ is drawn from a uniform distribution on the interval $[0, 2\pi]$.

The function ϕ_θ **inherits** randomness from the random variable θ.

A particular draw of the parameter, say $\theta = \pi$, then determines an instance of the neural network, here $\phi_\pi(t) = \cos(t + \pi)$.

In the picture of a conventional neural network before, the parameters are the weights $w_{ij}^{(n)}$ and biases $b_i^{(n)}$ at each layer of the network.

NN = $0d$ QFT.

Given a NN architecture ϕ_θ and parameter density $P(\theta)$, get correlators:

$$G^{(n)}(x_1, \ldots, x_n) = \langle \phi(x_1) \ldots \phi(x_n) \rangle = \int d\theta \, P(\theta) \, \phi_\theta(x_1) \ldots \phi_\theta(x_n) \, .$$

Viewing $P(\theta) = e^{-S[\theta]}$, this is a **0-dimensional quantum field theory**.

$$\begin{aligned} 0d \; &: \qquad \mathcal{Z}^{(0)} = \int_{\theta \in \mathbb{R}} d\theta \, e^{-S[\theta]} \, , \\ 1d \; &: \qquad \mathcal{Z}^{(1)} = \int_{\text{Paths } x(t)} \mathcal{D}x(t) \, e^{-S[x(t)]} \, , \\ d \geq 2 \; &: \qquad \mathcal{Z}^{(d)} = \int_{\text{Functions } \phi(x)} \mathcal{D}\phi(x) e^{-S[\phi(x)]} \, . \end{aligned}$$

From the $0d$ point of view, correlators $G^{(n)}$ are just observables built from the “fields” θ that depend on auxiliary variables x_i.

When do you "grow dimensions"?

Key question in NN-FT.

Under what conditions do the observables $G^{(n)}$ in the $0d$ QFT $(\phi_\theta, P(\theta))$ reproduce correlators of a QFT in $d > 0$ dimensions?

In particular, we wish to know when the $G^{(n)}$ satisfy the **Osterwalder - Schrader axioms** which allow reconstruction of a local, unitary Lorentzian QFT from the Euclidean correlation functions computed by the NN.

1. **Euclidean covariance**: the $G^{(n)}$ should transform covariantly under rotations and translations.
2. **Cluster decomposition**: correlators factorize, $G^{(p+q)} \to G^{(p)} G^{(q)}$, as the p points are taken far from the q points.
3. **Reflection positivity**: certain Euclidean correlators that compute norms $\langle \psi \mid \psi \rangle$ in the Lorentzian Hilbert space must be non-negative.

When do you "grow dimensions"?

Key question in NN-FT.

Under what conditions do the observables $G^{(n)}$ in the $0d$ QFT $(\phi_\theta, P(\theta))$ reproduce correlators of a QFT in $d > 0$ dimensions?

In particular, we wish to know when the $G^{(n)}$ satisfy the **Osterwalder - Schrader axioms** which allow reconstruction of a local, unitary Lorentzian QFT from the Euclidean correlation functions computed by the NN.

1. **Euclidean covariance**: the $G^{(n)}$ should transform covariantly under rotations and translations.
2. **Cluster decomposition**: correlators factorize, $G^{(p+q)} \to G^{(p)} G^{(q)}$, as the p points are taken far from the q points.
3. **Reflection positivity**: certain Euclidean correlators that compute norms $\langle \psi \mid \psi \rangle$ in the Lorentzian Hilbert space must be non-negative.

Example: Gaussian free field.

The standard construction of the free scalar in math literature defines

$$\varphi(x) = \sum_{k=1}^{\infty} \theta_k f_k(x) \, ,$$

where the f_k are an orthonormal basis of functions on some $\Omega \subset \mathbb{R}^d$ and the θ_k are i.i.d. random variables drawn from a unit normal distribution.

We interpret this as a neural network with infinitely many parameters θ_k.

This gives one example where a neural network ($0d$ field theory) can "mock up" the behavior of a higher-dimensional field theory.

Example: Gaussian free field.

The standard construction of the free scalar in math literature defines

$$\varphi(x) = \sum_{k=1}^{\infty} \theta_k f_k(x) \, ,$$

where the f_k are an orthonormal basis of functions on some $\Omega \subset \mathbb{R}^d$ and the θ_k are i.i.d. random variables drawn from a unit normal distribution.

We interpret this as a neural network with infinitely many parameters θ_k.

This gives one example where a neural network (0d field theory) can "mock up" the behavior of a higher-dimensional field theory.

NNs with infinitely many parameters.

The preceding example suggests that, in order to engineer NN-QFTs, it is important to have infinitely many parameters. What behaviors are possible in such infinite-parameter limits? There are at least three of interest:

1. Gaussian process limits. These involve an infinite-width limit which, by the CLT, leads to NNs with normally distributed outputs.
2. Stochastic process limits. NNs $\phi_\theta : \mathbb{R} \to \mathbb{R}$ written as sums of real-valued functions with random coefficients give stochastic processes, or theories of Euclidean quantum mechanics ($1d$ QFTs).
3. Distributional limits. NNs such as the Gaussian free field example in $d \geq 2$ diverge almost surely for any input x, but can be integrated against test functions. These are Euclidean QFTs in $d \geq 2$.

The initial work on NN-FT [Halverson, Maiti, Stoner '20] focused on Gaussian process limits. Let us briefly review the main idea of these limits.

NNs with infinitely many parameters.

The preceding example suggests that, in order to engineer NN-QFTs, it is important to have infinitely many parameters. What behaviors are possible in such infinite-parameter limits? There are at least three of interest:

1. Gaussian process limits. These involve an infinite-width limit which, by the CLT, leads to NNs with normally distributed outputs.
2. Stochastic process limits. NNs $\phi_\theta : \mathbb{R} \to \mathbb{R}$ written as sums of real-valued functions with random coefficients give stochastic processes, or theories of Euclidean quantum mechanics ($1d$ QFTs).
3. Distributional limits. NNs such as the Gaussian free field example in $d \geq 2$ diverge almost surely for any input x, but can be integrated against test functions. These are Euclidean QFTs in $d \geq 2$.

The initial work on NN-FT [Halverson, Maiti, Stoner '20] focused on Gaussian process limits. Let us briefly review the main idea of these limits.

The NN-GP correspondence.

We return to the single hidden layer architecture from before,

$$\phi_\theta(x) = \frac{1}{\sqrt{N}} \sum_{i=1}^{N} \sum_{j=1}^{d} w_i^{(1)} \sigma\left(w_{ij}^{(0)} x_j \right) ,$$

with biases set to zero for simplicity. The parameters in this NN are $\theta = \{ w_{ij}^{(0)}, w_i^{(1)} \}$ which we take independent and identically distributed.

Essentially due to the central limit theorem, we have the following.*

Theorem 2 (Neal)

As $N \to \infty$, the distribution of $\phi_\theta(x)$ approaches a Gaussian process:

$$\lim_{N \to \infty} \phi_\theta(x) \sim \mathcal{N}\left(\mu(x), K(x,y) \right) ,$$

with mean $\mu(x)$ and kernel $K(x,y)$ fixed by the distribution of weights.

*There is a nice visualization of this limit here.

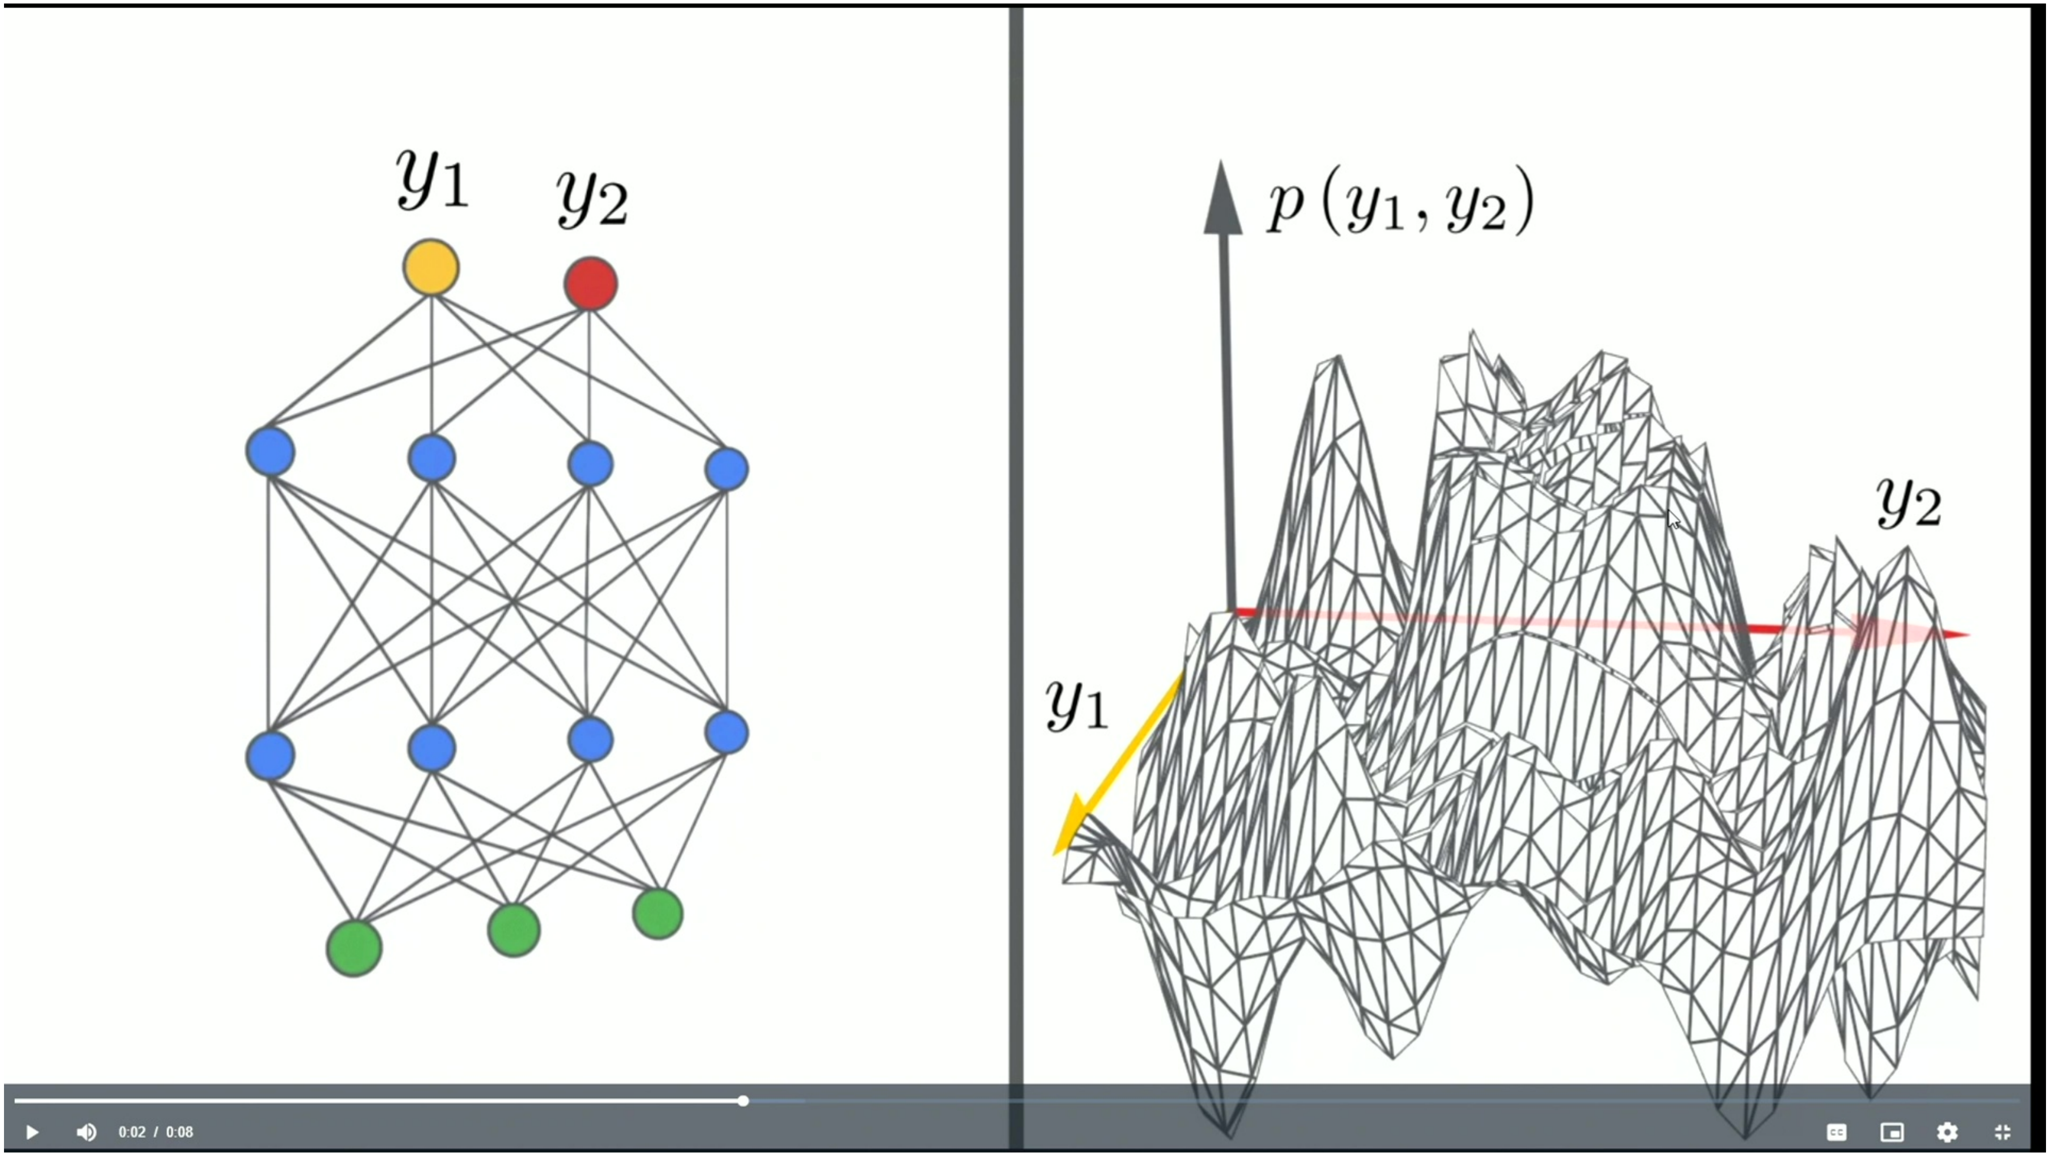

y_1
y_2
$p\left(y_1, y_2\right)$
y_2
y_1
0:02 / 0:08

A GP is *like* a free field theory...

A *Gaussian process* is a family of random functions $\phi_\theta(x)$ such that, for every finite set of inputs $x_1, \ldots, x_k$, the joint probability distribution

$$\mathbb{P}\left(\phi_\theta(x_1), \ldots, \phi_\theta(x_k)\right)$$

is a multivariate normal distribution; the characteristic function is

$$\mathbb{E}\left[\exp\left(i\sum_{\ell=1}^{k} a_k \phi_\theta(x_l)\right)\right] \stackrel{!}{=} \exp\left(-\frac{1}{2}\sum_{i,j} a_i a_j K(x_i, x_j) + i\sum_{\ell} a_\ell \mu(x_\ell)\right) .$$

If $\mu = 0$, this is like a source-free "free field theory" with associated action

$$S[\phi] = \int d^d x\, d^d y\, \phi(x) K(x,y)^{-1} \phi(y) \,,$$

where $\int d^d y\, K(x,y)^{-1} K(y,z) = \delta^{(d)}(x-z)$.

...but a GP is never *really* a QFT.

All moments of Gaussian distributions are finite for any values of the inputs. Therefore in any Gaussian process, one has

$$\lim_{y \to x} \langle \phi_\theta(x) \phi_\theta(y) \rangle < \infty \, .$$

However, for any unitary QFT in $d \geq 2$, the 2-point function of non-trivial local operators must diverge at coincident points.*

Thus, while the GP limit of NNs gives an interesting theory of random *functions*, it can never reproduce the correlators of a unitary QFT.

This is intuitively reasonable. We know that for $d \geq 2$, the path integral measure is supported on distributions rather than functions.

In $d = 1$, the path integral measure (for the free case, Wiener measure) is supported on continuous paths which are differentiable nowhere.

*One can argue for this by the Källén-Lehmann representation or by assuming the QFT is controlled by a unitary CFT in the UV, where $\langle \phi(x) \phi(y) \rangle = |x - y|^{-2\Delta}$.

Part 3: New results on neural network quantum mechanics.

QFT in one spacetime dimension.

Restricting to $d = 1$, a quantum mechanical theory is defined by correlators

$$G^{(n)}(t_1, \ldots, t_n) = \langle x(t_1) \ldots x(t_n) \rangle \, ,$$

which depend on n time coordinates $t_i \in \mathbb{R}$.

Such a theory is closely related to the notion of a *stochastic process*, by which we mean a collection of real-valued random variables

$$x_t = \{ x(t) \mid t \in T \} \, ,$$

where T is the *index set* for the process (either $\mathbb{R}$ or $[a, b]$).

The defining data of a stochastic process x_t is an assignment, for any collection of n variables $t_i \in T$ and for any n, of a joint PDF

$$\mathbb{P} \left(x(t_1), \ldots, x(t_n) \right) \, .$$

Requirements of quantum models.

Every stochastic process x_t defines a set of correlators $G^{(n)}$. But not every collection of $G^{(n)}$ is suitable for defining a quantum theory.

Eventually we wish to impose Osterwalder-Schrader. But first we work with "minimal quantum models" or MQMs which only obey:

1. Mean-square continuity: $\lim_{t\to s} \left\langle |x(t) - x(s)|^2 \right\rangle = 0$.
2. The $1d$ version of Källén-Lehmann:

$$\langle x(t)x(s)\rangle = \int_0^\infty dm\, \rho(m) e^{-m|t-s|} \overset{\text{e.g.}}{=} \sum_{n=1}^{\infty} e^{-E_n|t-s|} \left|\langle 0 \mid x \mid n\rangle\right|^2 ,$$

where $\rho(m) \overset{\text{e.g.}}{=} \sum_{n=1}^{\infty} \left|\langle 0 \mid x \mid n\rangle\right|^2 \delta\left(m - E_n\right)$ is normalizable.

The first condition implies that sample paths $x(t)$ must be continuous, and that $G^{(2)}$ is continuous. The second implies that $G^{(2)}$ is always finite.

Requirements of quantum models.

Every stochastic process x_t defines a set of correlators $G^{(n)}$. But not every collection of $G^{(n)}$ is suitable for defining a quantum theory.

Eventually we wish to impose Osterwalder-Schrader. But first we work with "minimal quantum models" or MQMs which only obey:

1. Mean-square continuity: $\lim_{t \to s} \left\langle |x(t) - x(s)|^2 \right\rangle = 0$.
2. The 1d version of Källén-Lehmann:

$$\langle x(t) x(s) \rangle = \int_0^\infty dm \, \rho(m) e^{-m|t-s|} \overset{\text{e.g.}}{=} \sum_{n=1}^{\infty} e^{-E_n |t-s|} \left| \langle 0 \mid x \mid n \rangle \right|^2 ,$$

where $\rho(m) \overset{\text{e.g.}}{=} \sum_{n=1}^{\infty} \left| \langle 0 \mid x \mid n \rangle \right|^2 \delta(m - E_n)$ is normalizable.

The first condition implies that sample paths $x(t)$ must be continuous, and that $G^{(2)}$ is continuous. The second implies that $G^{(2)}$ is always finite.

A universality theorem for NN-QM.

One of the results of our paper is the following.

Theorem 3

Any Euclidean quantum mechanical model $x(t)$ satisfying assumptions (1) and (2) admits a neural network description $x_t = \phi_\theta(t)$ with architecture

$$\phi_\theta(t) = \langle x(t) \rangle + \sum_{k=1}^{\infty} \theta_k f_k(t) \,,$$

where $f_k(t)$ are continuous, orthogonal real-valued functions and θ_k are pairwise uncorrelated random variables.

The proof of this result relies on the Kosambi-Karhunen-Loève theorem from the literature on stochastic processes.

This means that every MQM is a NN-SP, a stochastic process which admits a NN description. We also use the term NN-QM for such models.

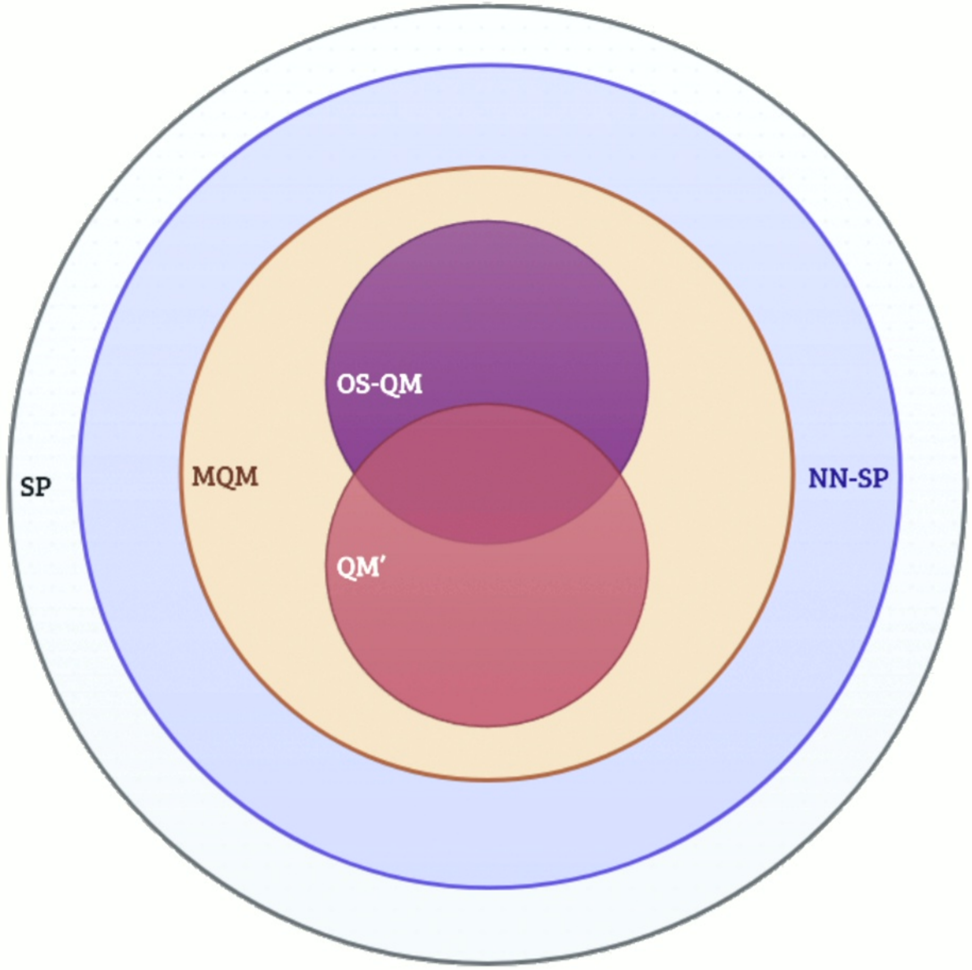

OS-QM
SP
MQM
NN-SP
QM′

OS axioms.

Under what conditions can such a Euclidean quantum model be analytically continued to real time? We need to satisfy the OS axioms:

1. Euclidean covariance: the stochastic process x_{t+a} should be equivalent* to x_t, for any a, and the process x_{-t} should be equivalent to x_t. That is, x_t should be a *stationary symmetric process*.

2. Cluster decomposition: the stochastic process x_t should be *mixing*, which implies ergodicity.

3. Reflection positivity: for any collection of positive times $t_i > 0$ and any bounded function F of n variables,

$$\langle F(x(t_1), \ldots, x(t_n)) \left(F(x(-t_1), \ldots, x(-t_n)) \right)^* \rangle \geq 0 \,.$$

*Two stochastic processes x_t and y_t are said to be *equivalent* if, for any collection of times $t_1, \ldots, t_n$, the joint PDFs obey $\mathbb{P}(x(t_1), \ldots, x(t_n)) = \mathbb{P}(y(t_1), \ldots, y(t_n))$.

Mechanisms for reflection positivity.

The reflection positivity (RP) condition is interesting and subtle. In our paper, we explore two ways to engineer RP for NNs or SPs:

1. A "parameter-splitting" mechanism for NN-QM, where we assume $\{\theta\} = \{\theta^0\} \cup \{\theta^+\} \cup \{\theta^-\}$ where $\phi_\theta(t)$ is independent of θ^+ for $t < 0$ and independent of θ^- for $t > 0$. This lets us write

$$\begin{aligned}\langle F\,(TF)^*\rangle &= \int d\theta^0\, P(\theta^0) \left(\int d\theta^+\, P_+(\theta^0, \theta^+)\, F\left(\phi_\theta(t_1), \ldots, \phi_\theta(t_n)\right) \right) \\ &\quad \cdot \left(\int d\theta^-\, P_-(\theta^0, \theta^-)\, \left(F\left(\phi_\theta(-t_1), \ldots, \phi_\theta(-t_n)\right)\right) \right)^* \\ &= \int d\theta^0\, P(\theta^0) \left| \int d\theta^+\, P_+(\theta^0, \theta^+)\, F\left(\phi_\theta(t_1), \ldots, \phi_\theta(t_n)\right) \right|^2 .\end{aligned}$$

2. If x_t is a symmetric *Markov* process, then it is RP. This follows from work of Nelson and OS, but we present an elementary proof for $1d$.

Markov is fragile but RP is robust.

If x_t and y_t are Markov processes, then $z_t = x_t + y_t$ need not be Markov, and $f(x_t)$ for a deterministic function f need not be Markov.

However, linear combinations of RP processes are RP. We prove more:

Theorem 4

Let x_t be a reflection-positive stochastic process and consider a family of bounded random functions realized by a neural network architecture ϕ_θ with parameter density $P(\theta)$. Then the stochastic process

$$y_t = \phi_\theta(x_t)$$

is also reflection-positive.

This facilitates the construction of *deep NN-QM models*. Choose a collection of RP processes $x_t^{(i)}$ (e.g. symmetric Markov) as inputs to a NN with random parameters. Then the output y_t is also RP.

Spectrum of example deep NN-QM model.

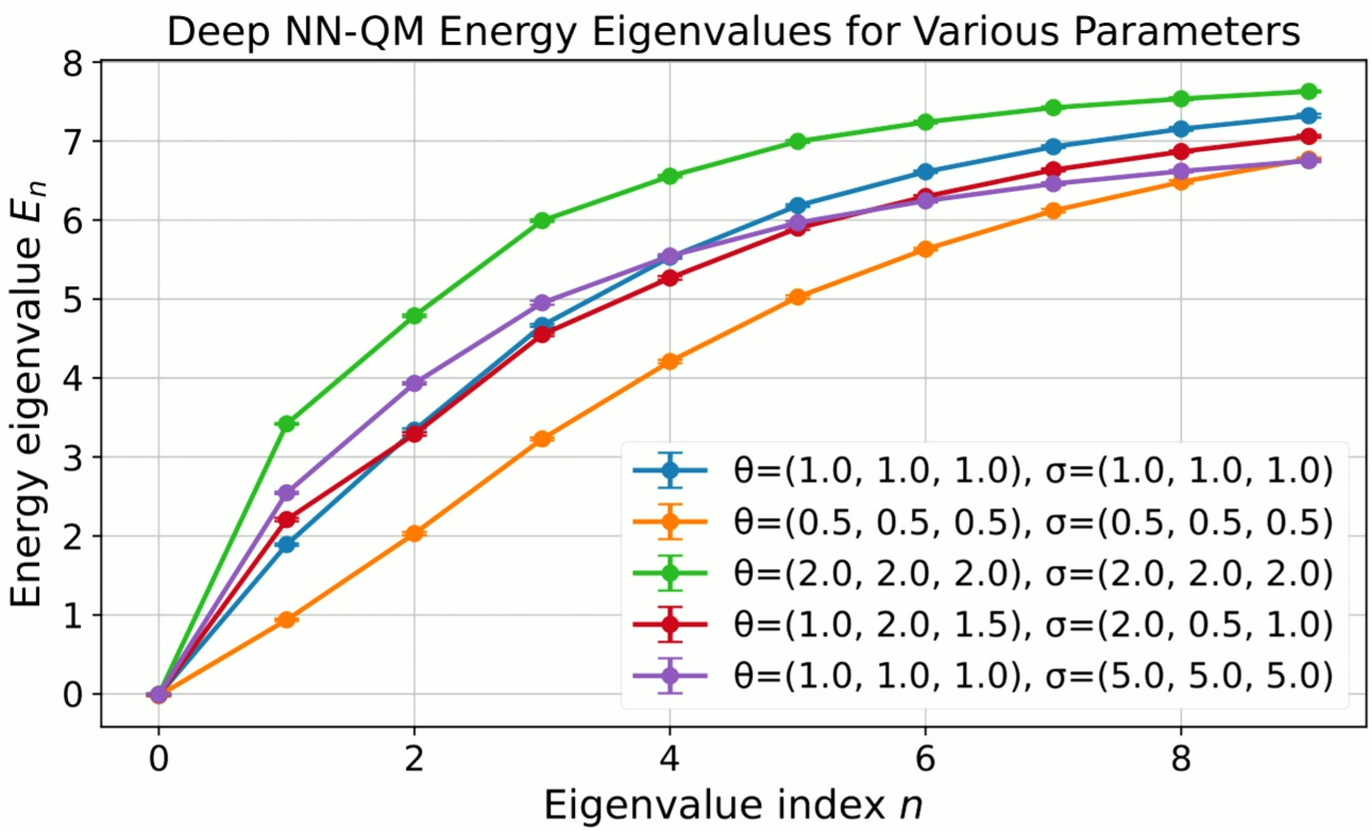

Part 4: Summary and future directions.

Summary.

We have given an overview of applications of neural networks to quantum field theory, focusing on the $1d$ (QM) case where one has greater control.

Some take-away messages include:

1. Neural networks ϕ_θ with random parameters θ give a powerful way of parameterizing random functions. Often there are universality results.
2. To reproduce field theory correlators, we need NNs with infinitely many parameters. Depending on how such infinite-parameter limits are taken, different behavior can result (GP, SP, distributional NN...).
3. Under mild assumptions, any $1d$ field theory admits a NN description.
4. OS axioms can be engineered for $1d$ NN-QM systems in several ways.
5. Applying any NN to a RP quantum system preserves RP. One can construct a large class of unitary quantum models using deep NNs.

Further research.

There are several potential directions for extending our analysis.

1. Use *training* to construct NN-QM models that reproduce desired features (correlators, energies, etc.) of a "target" QM theory.
2. Prove or disprove the conjecture that *every* reflection-positive process can be realized as a deep NN-QM with symmetric Markov inputs.
3. Construct interacting $d \geq 2$ field theories using neural networks, which requires understanding NN distributions (generalized functions).

Progress on these directions could improve our understanding of functional integration using neural network tools, and perhaps help us develop a clearer picture of what a quantum field theory really is.

Thank you for your attention!