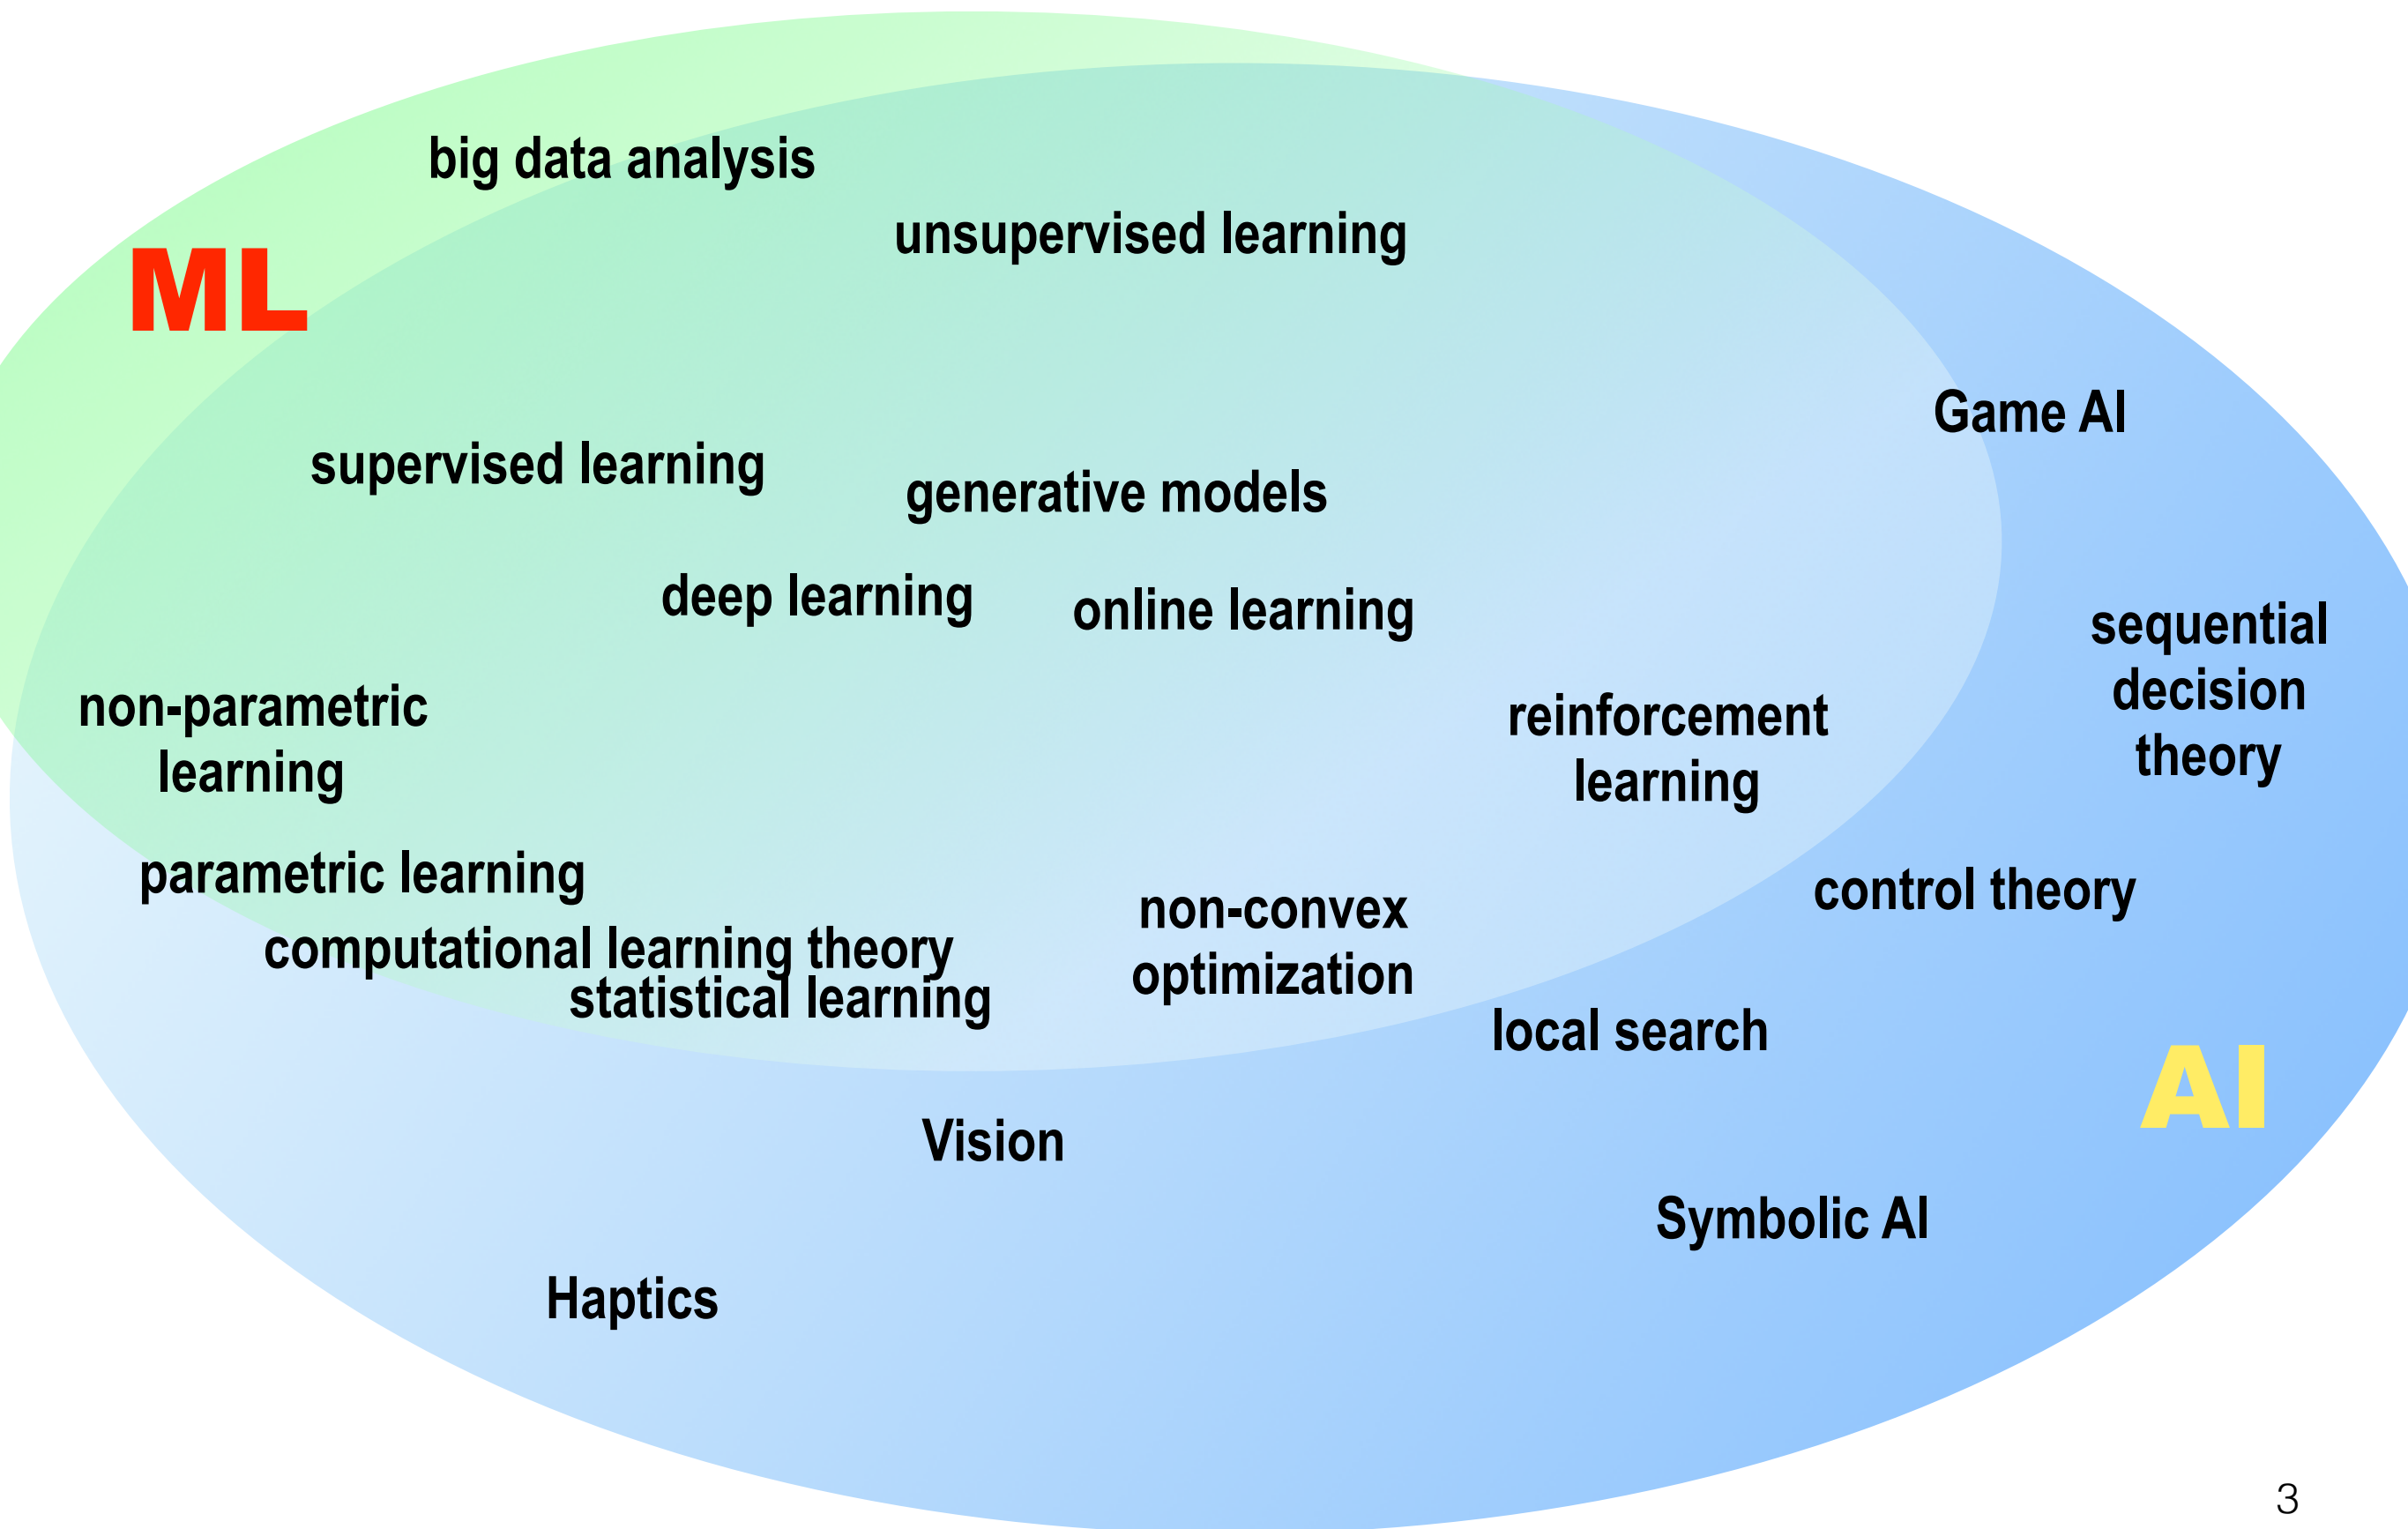


Quantum-enhanced Machine Learning (with near-term devices)

Contents & Literature

- 1) *Background 1: machine learning (ML)*
 - *what is ML, and basic ML models*
- 2) *QC meets ML (big picture) [for more info: arXiv:1709.02779]*
- 3) *ML and parametrized circuits [for more info: arXiv:1906.07682]*
- 4) *QeML with quantum feature spaces [based on: arXiv:1804.11326]*
 - *Support vector machines*
 - *Explicit and implicit quantum-embedded SVMs*

Machine learning and AI

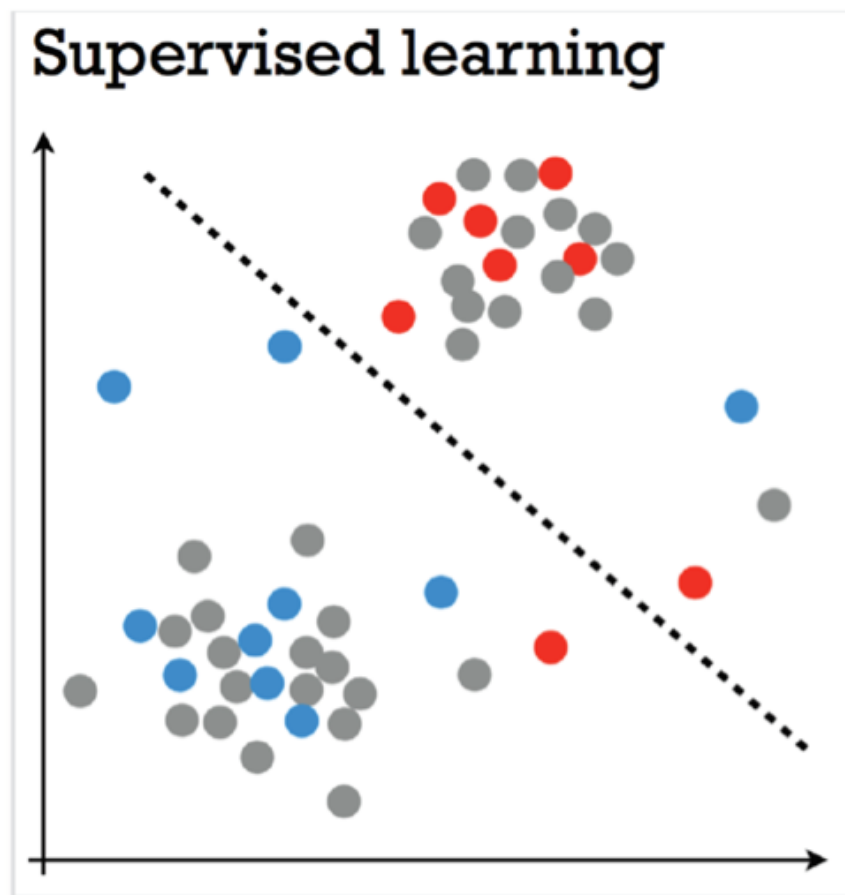


Three main (cannonical) modes of ML:

- Supervised learning
- Unsupervised learning
- Reinforcement learning

- forest of in-between models; semi-supervised, active, transductive, on-line...

Supervised learning: the *what* (is the objective)



?

Supervised learning: the *what* (is the objective)

Basic concepts and math

Data (feature vectors) & Labels:

$$\mathbf{x} \in S \subseteq \mathbb{R}^n; y \in Labels$$

Label function:

$$f : S \rightarrow Labels$$

Dataset, “training examples”

$$D = \{(\mathbf{x}_i, y_i) \mid \mathbf{x}_i \in S; y_i = f(\mathbf{x}_i)\}$$

-need to correctly label unlabeled data

Given D , output a good guess for f .

-*classification (categorical or discrete label) v.s. regression (continuous label)*

-*classification, prediction, regression....*

Supervised learning: the *what* (is the objective)

Basic concepts and math

More generally (probabilistic)

BTW: Distributions generalize functions

Data (feature vectors) & Labels:

$$\mathbf{x} \in S \subseteq \mathbb{R}^n; y \in Labels$$

Label function:

$$P(\mathbf{x}, y)$$

Dataset, “training examples”

$$D \sim P^{\times |D|}$$

Given D , output a good guess for $P(y | \mathbf{x})$

Learning about data-label relationships in a bivariate distribution from samples

Unsupervised learning: the *what* (is the objective)



?

data: $\mathbf{x} \in S \subseteq \mathbb{R}^n$

“world”: $P(\mathbf{x})$

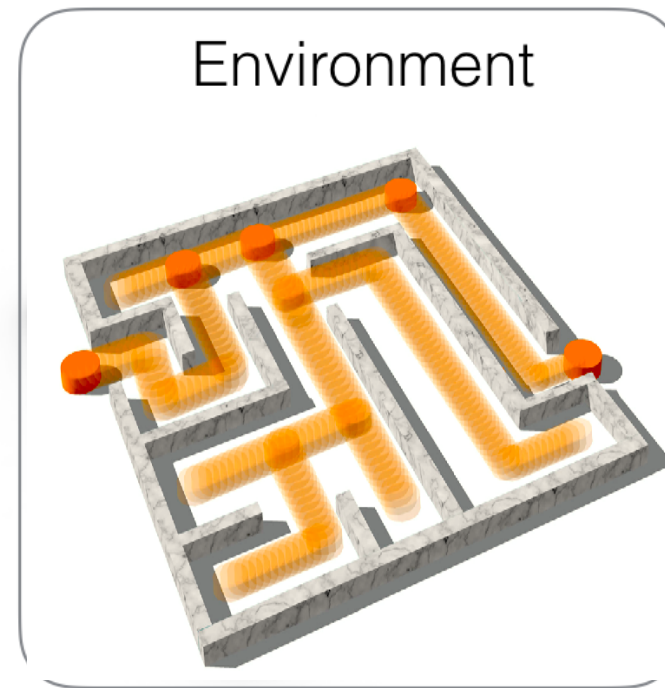
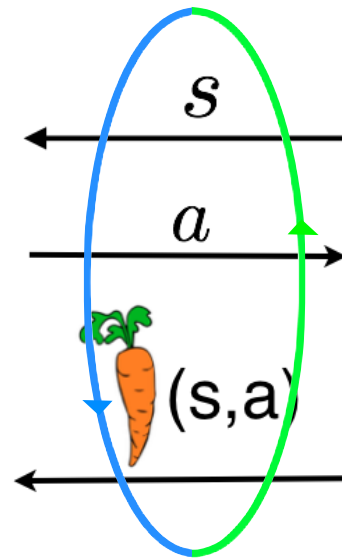
training: $D \sim P^{\times |D|}$

-discriminative (clustering) , “labeling w/o examples”

-generative (make more cats):
approximate sampling from P given D

Learning about (all) features in a distribution from samples

Reinforcement learning: the *what* (is the objective)



$$\mathcal{S} = \{s_1, s_2, \dots\}$$
$$\mathcal{A} = \{a_1, a_2, \dots\}$$

$$\pi(a|s)$$

$$T(s|s', a)$$

Learning correct behaviour (policies) by trial-and-error
(incl. data generation online). E.g. AlphaGo.

Supervised learning: the *how* (is it achieved)

Recall: need to “guess” $f : S \subseteq \mathbb{R}^n \rightarrow \text{Labels}$ from $D = \{(\mathbf{x}_i, y_i = f(\mathbf{x}_i))\}$

- Hypothesis family: $\{f^\theta \mid f^\theta : S \subseteq \mathbb{R}^n \rightarrow \text{Labels}, \theta\}$ (c.f. “**model/model family**”)

- Learning = training \approx fitting:

$$\operatorname{argmin}_\theta \text{Error_on_}D(f^\theta) + R(f^\theta)$$

$R = \text{regularization}$

- “Loss”, “empirical risk”, e.g. $\sum_{(\mathbf{x}, y) \in D} |f^\theta(\mathbf{x}) - y|^2$

- Generalization performance: (no overfitting, Occam’s razor)

Supervised learning: the *how* (is it achieved)

Recall: need to “guess” $f : S \subseteq \mathbb{R}^n \rightarrow \text{Labels}$ from $D = \{(\mathbf{x}_i, y_i = f(\mathbf{x}_i))\}$

- Hypothesis family: $\{f^\theta \mid f^\theta : S \subseteq \mathbb{R}^n \rightarrow \text{Labels}, \theta\}$ (c.f. “**model/model family**”)

- Learning = training \approx fitting:

$$\operatorname{argmin}_\theta \text{Error_on_}D(f^\theta) + R(f^\theta)$$

$R = \text{regularization}$


the same elements will
be present for unsupervised
learning

- “Loss”, “empirical risk”, e.g. $\sum_{(\mathbf{x}, y) \in D} |f^\theta(\mathbf{x}) - y|^2$


- Generalization performance: (no overfitting, Occam’s razor)

Regularization: controlling "model complexity" to ensure good generalization

Classification



Regression




Supervised learning: the *how* (is it achieved); examples

- support vector machines (SVM)
- neural networks

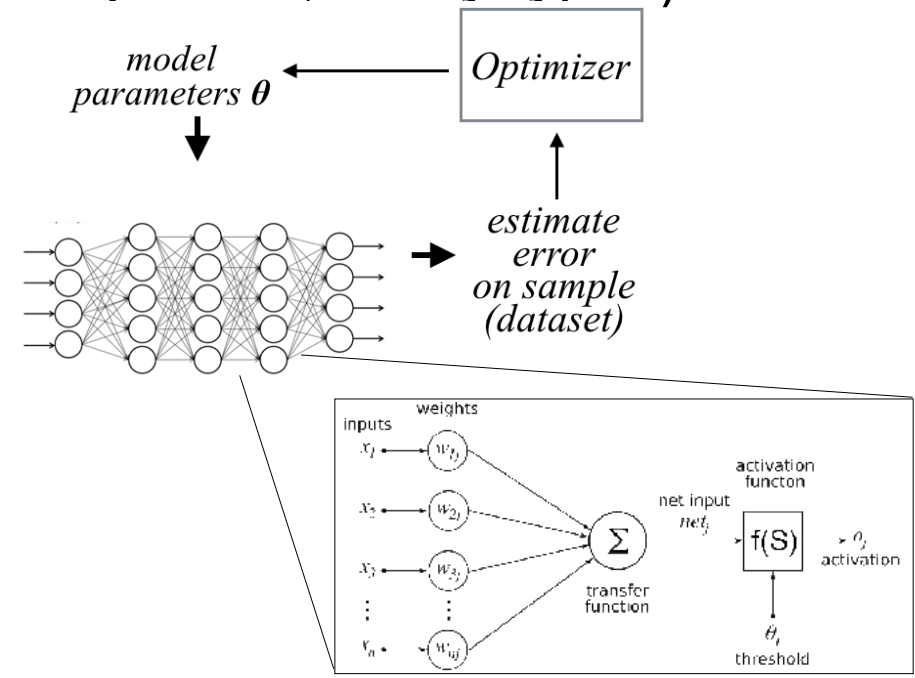
- k-nearest neighbours [classification]
- decision trees [classification]
- naïve Bayes
- (linear) regression [regression]
- Gaussian process regression

Supervised learning: the *how* (is it achieved); examples

Support Vector Machines (SVMs)



Neural networks (not just SL; many types)



“which half-space” \forall hyperplanes
 $sign(\mathbf{w}^t \cdot \mathbf{x} + b)$ + important trick

hypothesis family

specified by n.n.
 (linear+nonlinear layers)

normal vector+offset

model parameters

weights + offsets

usually Lagrangian multiplier method

training/ optimization

usually backpropagation
 (chain-rule based (stochastic) gradient descent)

regularization?

When discussing QML, keep an eye on

- the **What**
(what is the objective/goal)
- the **How**
(how is it done: algorithm; does it achieve the goal)
- the **Why**
(why do it on a QC; what is the expected advantage/other motivation)

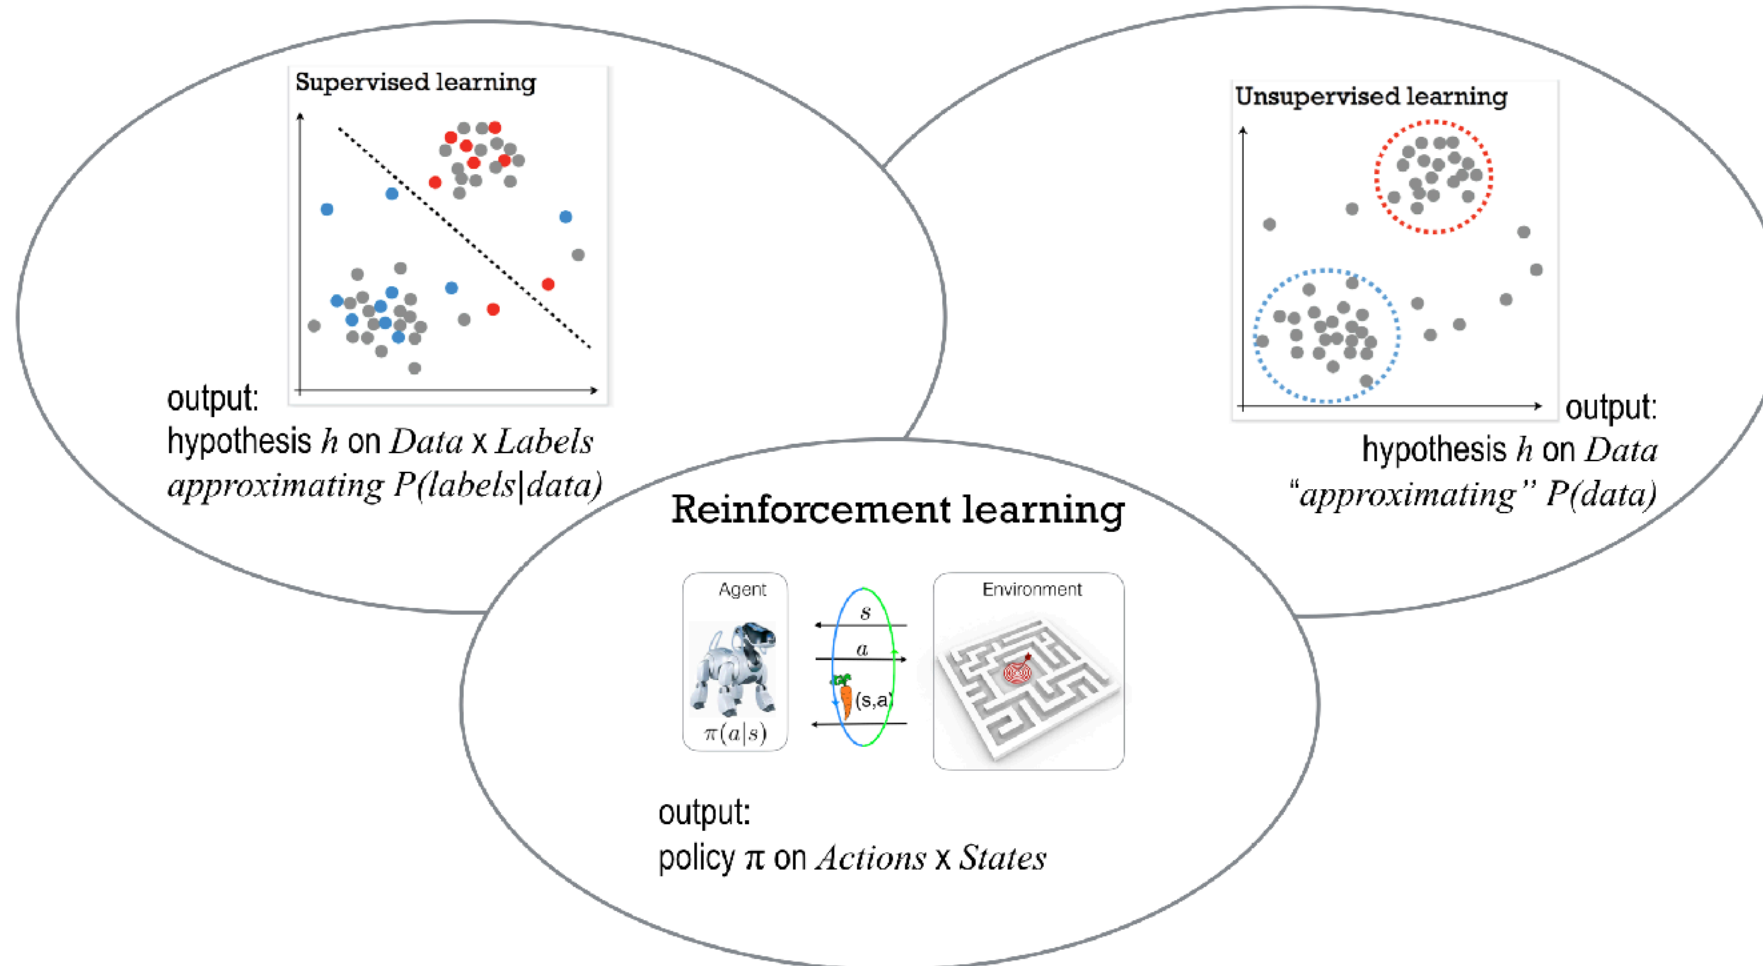
actually, same questions apply to much of *classical* ML approaches

the why is tricky tho; makes a good model model is though

Big picture take home:

the learning/training is optimization:

$$\operatorname{argmin}_{\theta} \operatorname{Err}_{\text{training_set}}(\theta) + \operatorname{Reg}(\theta)$$



but machine learning is more; which model; how it generalizes; good choices...

A connection...

variational methods in physics.. incl VQE are very similar

| Var. Q chem | ML |
|-----------------------------------|---------------------------------|
| "Ansatz" | "model"/hypothesis family |
| loss: energy | loss: training set error+regul. |
| explicit, error free ground truth | implicit ground truth, errors |
| optimization | learning/training |



no regularization
or generalization

statistical,
parameteric learning

regularization
generalization

-ground truth...

Cat v.s. no-cat example



QC meets ML: big picture ideas

-QC and the optimization bottleneck

training is optimization and can be hard (NP-hard) → quantum optimization

-QC and the high dimension bottleneck

much of ML is linear algebra; quantum computing is good at that, under conditions

-QC and the hard model bottleneck

topic of the next two lectures

Supervised Machine learning with Parameterized Quantum Circuits

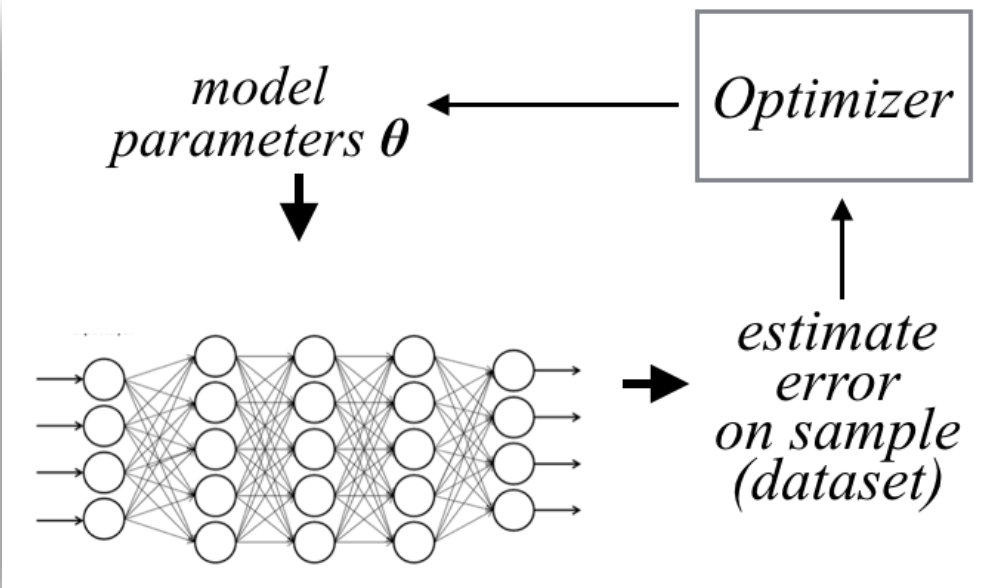
What: supervised learning for classification

Using quantum computing... but not for optimization needs

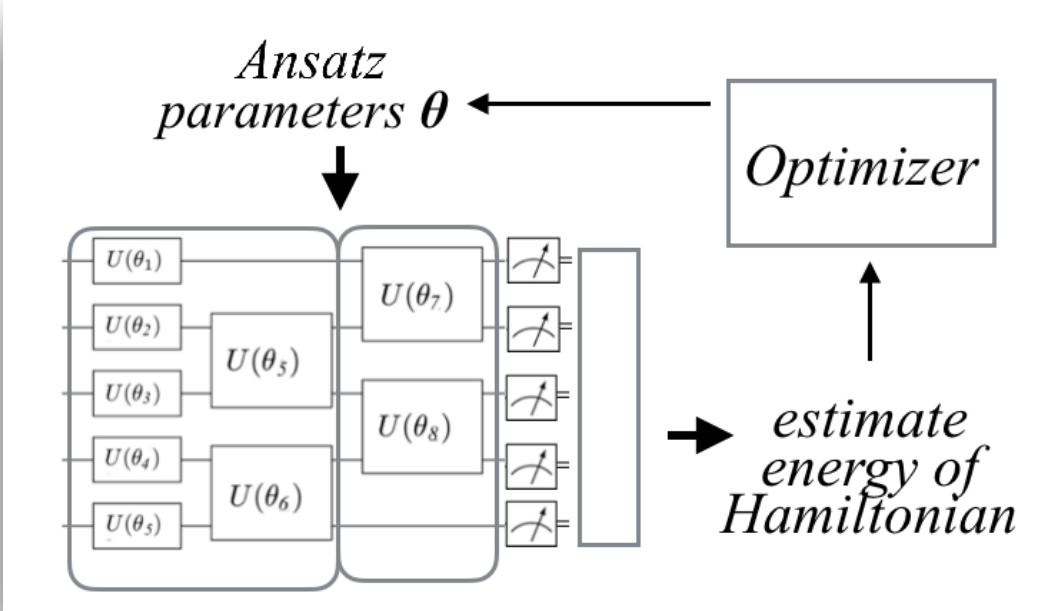
Why? TBD

Supervised Machine learning with Parameterized Quantum Circuits

neural networks

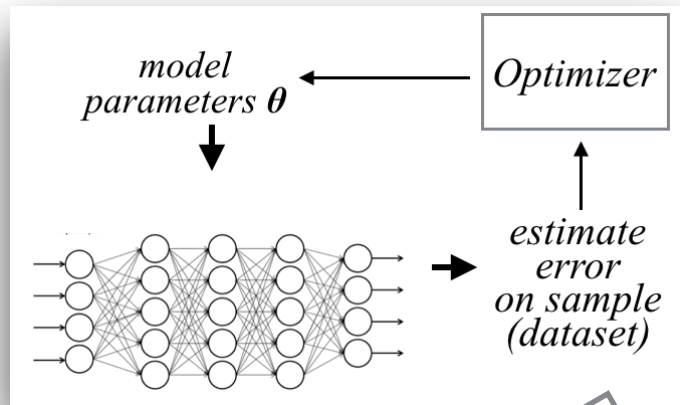


VQE

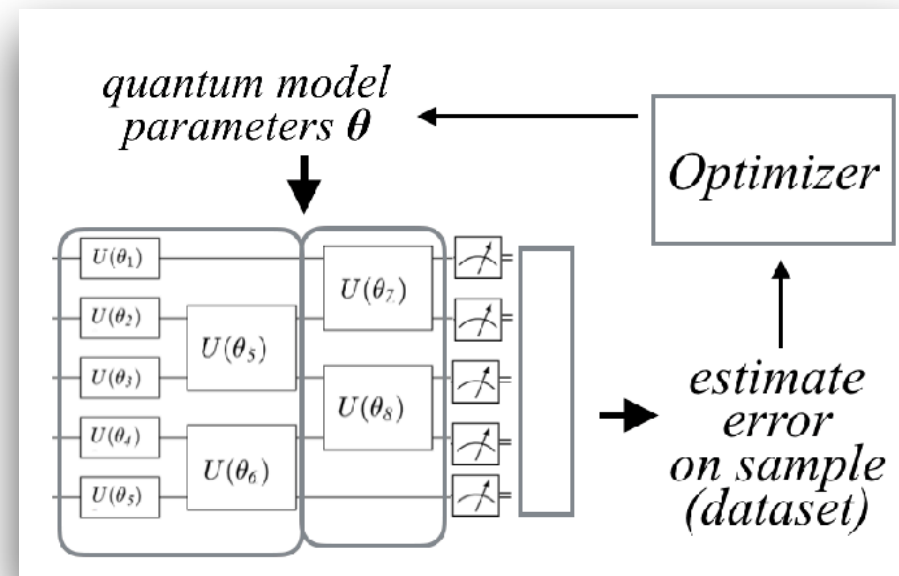
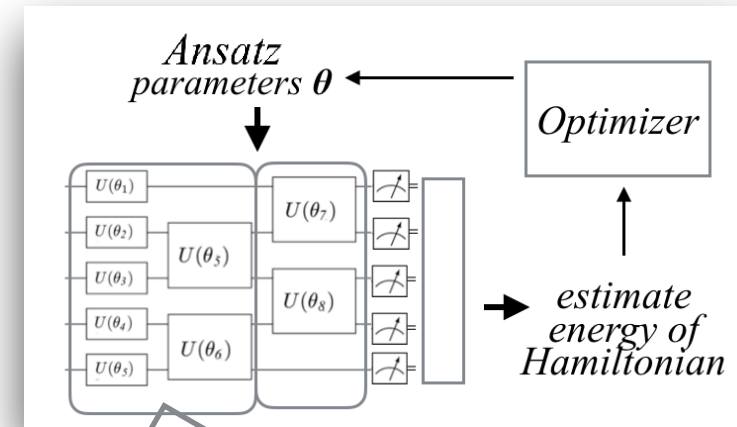


Machine learning with Parameterized Quantum Circuits

neural networks



VQE



PQC-based ML

- 1) can we train it?
- 2) does it work?
- 3) does it do anything interesting?
why do this?

Same Q's for VQE, but there 3) is clear. Here it is not.

Motivations: cannot do it classically? Curiosity driven?


We don't really understand the model...

Next:



- 1) a way to understand some of it.
- 2) reasons to do it

Background 2: SVMs in detail

$$D = \{(x_i, y_i)\}_i \quad x_i \in \mathbb{R}^d, y_i \in \{-1, 1\}$$



*separating hyperplanes
(linear classifier, not SVM)*



SVM: max-margin hyperplanes

for now, assume linearly separable data


$$D = \{(x_i, y_i)\}_i \quad x_i \in \mathbb{R}^d, y_i \in \{-1, 1\}$$

A number of equivalent formulations...

$$\arg \max_{\mathbf{w}, b} \min_{i \in \{1, \dots, N\}} \frac{y_i(\mathbf{w}^\top \mathbf{x}_i + b)}{\|\mathbf{w}\|}$$

$y_i(\mathbf{w}^\top \mathbf{x}_i + b)$ - “functional margin”

$\frac{y_i(\mathbf{w}^\top \mathbf{x}_i + b)}{\|\mathbf{w}\|}$ - “geometric margin”




$$D = \{(x_i, y_i)\}_i \quad x_i \in \mathbb{R}^d, y_i \in \{-1, 1\}$$

After some work: quadratic problem

$$\arg \min_{\mathbf{w}, b} \frac{1}{2} \|\mathbf{w}\|^2$$

such that $y_i(\mathbf{w}^\top \mathbf{x}_i + b) \geq 1, \quad i = 1, \dots, N.$



SVM: max-margin hyperplanes

“*Support vectors*”:
points closest and equidistant
to hyperplane


Hyperplane fully defined
in terms of *support vectors*

Lagrangian approach

Primal problem:

$$\arg \min_{\mathbf{w}, b} \frac{1}{2} \|\mathbf{w}\|^2$$

such that $y_i(\mathbf{w}^T \mathbf{x}_i + b) \geq 1, \quad i = 1, \dots, N.$



Dual problem:

$$\arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j (\mathbf{x}_i)^T \mathbf{x}_j,$$

such that $\alpha_i \geq 0, \quad \text{for } i = 0, \dots, N,$

and $\sum_{i=1}^N \alpha_i y_i = 0.$

$$\mathbf{w} = \sum_{i=1}^N \alpha_i^* y_i \mathbf{x}_i.$$

Why bother with dual problem? Representation in *terms of datapoints*

- sparser evaluation
- only inner products matter
- handy for *quantum tricks*


$$(\mathbf{w}^*)^T \mathbf{x} + b^* = \left(\sum_{i=1}^N \alpha_i y_i (\mathbf{x}_i)^T \mathbf{x} \right) + b^*.$$

$$\alpha_i \alpha_j y_i y_j (\mathbf{x}_i)^T \mathbf{x}_j,$$

Comment: the math should not hide the fact we are simply
finding a member of the hypothesis family
which is minimizing a loss function

NB: almost true: SVMs is “optimized”
to be able to reason about learning performance...


*Why should we care about SVMs:
what about when data is **not** linearly separable?*



Non-separable datasets?

-slack variables (this lead to QSVM - type 1)

-feature mapping and the kernel trick



$$\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^D \quad \vec{x} \xrightarrow{\Phi} \Phi(\vec{x})$$

c.f.: Cover's theorem...

The kernel trick:


one can “train” and evaluate SVM classifiers in rich feature spaces without ever mapping data-points into said spaces. They can even be infinite dimensional

The kernel trick

Note: in dual... only inner products matter

$$K(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \quad (\phi = \Phi \dots)$$

$$\arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j (\mathbf{x}_i)^T \mathbf{x}_j,$$




$$\begin{aligned} & \arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \\ &= \arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j K(\mathbf{x}_i, \mathbf{x}_j) \end{aligned}$$

c.f. Mercer's theorem

The kernel trick

Note: in dual... only inner products matter

$$\arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j (\mathbf{x}_i)^{\top} \mathbf{x}_j,$$



$$K(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \quad (\phi = \Phi \dots)$$

BTW: this thing is called "the kernel" $\langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$

Note, we really don't care about the feature map Φ itself...

The kernel trick

Kernels can sometimes be evaluated (much) more efficiently directly:

E.g. (stupidly)

$$(x_1, x_2, x_3) \mapsto \phi(\mathbf{x}) = (x_1x_1 \quad x_1x_2 \quad x_1x_3 \quad x_2x_1 \quad x_2x_2 \quad x_2x_3 \quad x_3x_1 \quad x_3x_2 \quad x_3x_3)^T$$

$$\langle \phi(\mathbf{x}), \phi(\mathbf{z}) \rangle = \sum_{i=1}^d \sum_{j=1}^d x_i z_i x_j z_j \quad \text{Runtime for } \phi(\mathbf{x}): \mathcal{O}(d^2)$$

The kernel trick

$$\phi(\mathbf{x}) = (x_1x_1 \quad x_1x_2 \quad x_1x_3 \quad x_2x_1 \quad x_2x_2 \quad x_2x_3 \quad x_3x_1 \quad x_3x_2 \quad x_3x_3)^\top$$

reverse-engineered:
$$K(\mathbf{x}, \mathbf{z}) = (\mathbf{x}^\top \mathbf{z})^2 = \left(\sum_{i=1}^d x_i z_i \right) \left(\sum_{i=1}^d x_i z_i \right) = \sum_{i=1}^d \sum_{j=1}^d x_i z_i x_j z_j = \langle \phi(\mathbf{x}), \phi(\mathbf{z}) \rangle.$$

Directly:

Let $\mathbf{x} = (x_1, \dots, x_d)^\top$, $\mathbf{z} = (z_1, \dots, z_d)^\top$ and

$$K(\mathbf{x}, \mathbf{z}) = (\mathbf{x}^\top \mathbf{z})^2.$$

Runtime: $\mathcal{O}(d)$.

Yay, quadratic speedup

See e.g. *Radial basis function kernel*

$$K(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\sigma^2}\right)$$

$$K(x, x') = \langle \Phi(x), \Phi(x') \rangle$$

$$\Phi(x) = e^{-x^2/2\sigma^2} \left[1, \sqrt{\frac{1}{1!\sigma^2}} x, \sqrt{\frac{1}{2!\sigma^4}} x^2, \sqrt{\frac{1}{3!\sigma^6}} x^3, \dots \right]^\top$$

inf. dim.....

c.f. Mercer's theorem

To keep in mind:

-primal v.s. dual:

in primal, optimize over normal vector **explicitly**;

in dual, **it is implicit**, and the separating hyperplane is expressed in terms of data points

-feature maps:


by raising dimension non-linearly, we can achieve linear-separability

-kernel trick:

in dual formulation, only need kernel evaluation on data points for training.

Back to Quantum: an SVM reading of PQC-powered ML

Basic idea: quantum computing offers interesting “natively quantum” feature maps and kernels




$$\vec{x} \mapsto \mathcal{U}_{\Phi}(\vec{x})|0\rangle = |\Phi(\vec{x})\rangle$$

*Data is encoded **in the circuit parameters** (not input state). More general.*

Basic idea: quantum computing offers interesting “natively quantum” feature maps and kernels

One thing we can do with this... is evaluate inner products.



Kernel!

$$|\langle \Phi(\vec{y}) | \Phi(\vec{x}) \rangle|^2$$

Can be hard to compute.

Do this quantumly

(recall QC is good for inner products)

also possible:

swap tests, hadamard tests

But we can do more


Which feature maps should we construct?

$$U_{\Phi(\vec{x})} = \exp \left(i \sum_{S \subseteq [n]} \phi_S(\vec{x}) \prod_{i \in S} Z_i \right)$$

$$\phi_{\{i\}}(\vec{x}) = x_i \text{ and } \phi_{\{1,2\}}(\vec{x}) = (\pi - x_1)(\pi - x_2)$$

$$e^{i\phi_{\{l,m\}}(\vec{x})} Z_l Z_m = \begin{array}{c} \bullet \text{---} \text{---} \bullet \\ | \qquad \qquad | \\ \oplus \text{---} \boxed{Z_\phi} \text{---} \oplus \end{array}$$

$$\mathcal{U}_\Phi = H^{\otimes n} U_\Phi H^{\otimes n} U_\Phi \cdots H^{\otimes n} U_\Phi$$



Which feature maps should we construct...elaborated

- DIMENSION OF FEATURE SPACE = $2^{\# \text{QUBITS}}$; # QUBITS = N = INITIAL DIMENSION

- DEFINE "SUBMAPS" ϕ_S
 - S = - individual vector entries ; $S \subseteq \{1 \dots N\}$
 - pairs ; ... can be generalized
 - $:=$ correlators (2 OR k -local)

$$\phi_S : \mathbb{R} \text{ or } \mathbb{R}^2 \rightarrow \text{"angles"}$$

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix} \xrightarrow{\phi_S} \theta \rightarrow \underbrace{\exp(i z_1 \otimes z_3 \theta)}_{U_S(\vec{x})}$$

- $U_{\Phi}(\vec{x}) := \prod_S U_S(\vec{x}) \dots$ ALL DIAGONAL...

- FEATURE MAP : $\mathcal{U}_{\Phi} = \left(H^{\otimes N} U_{\Phi}(\vec{x}) \right)^{\otimes m}$ m- hyperparameter

First type of PQC SVM: implicit (dual) model

training:


$$\arg \max_{\alpha} \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j K(\mathbf{x}_i, \mathbf{x}_j)$$

classifying:

$$out(\mathbf{x}) = \text{sign} \left(\sum_{i=1}^N y_i \alpha_i^* K(\mathbf{x}_i, \mathbf{x}) + b \right)$$


Quantum parts: needed in both to evaluate the kernels

only offline; optimization essentially on classical data.




$$O(N^2 / \text{poly}(\epsilon))$$

Fully quantum model: explicit (primal) model



$f(z)$

$$f(z) : \{0, 1\}^n \rightarrow \{-1, 1\}$$



repeat l - times

$$U(\theta_{m,t}) = e^{i\frac{1}{2}\theta_{m,t}^z Z_m} e^{i\frac{1}{2}\theta_{m,t}^y Y_m}$$

How does it output a label? What is the achieved classifier?


The label (output): (approximately) sign of the expected value f , shifted by b :

$$out(\mathbf{x}, \theta) \approx \text{sign}(\mathbb{E}_{z \sim Q.circ}[f(z)] + b)$$

$$out(\mathbf{x}, \theta) \approx \text{sign}(\langle \Phi(\vec{x}) | W^\dagger(\vec{\theta}) \mathbf{f} W(\vec{\theta}) | \Phi(\vec{x}) \rangle + b)$$

The algorithm:

- sample z many times (“shots”)
- average, shift, compute sign.



Comment:

$$\text{out}(\mathbf{x}, \theta) \approx \text{sign}(\mathbb{E}_{z \sim Q.\text{circ}}[f(z)] + b)$$

$$\text{out}(\mathbf{x}, \theta) \approx \text{sign}(\langle \Phi(\vec{x}) | W^\dagger(\vec{\theta}) \mathbf{f} W(\vec{\theta}) | \Phi(\vec{x}) \rangle + b)$$

"measure each qubit in comp basis, compute $f(\vec{z})$ " = observable

$$\mathbf{f} := \sum_{\vec{z}} f(\vec{z}) |\vec{z}\rangle\langle\vec{z}|$$

"measure each qubit in comp basis, compute $f(\vec{z})$ " := a realization of measurement of \mathbf{f} . AVERAGING YIELDS THE EXPECTED VALUE.

How does it learn?

Optimize θ to minimize some loss/error/empirical risk on dataset

Involves evaluation of classifier function many times...

Often: stochastic gradient descent

Q. chemistry optimization and optimization here very similar

But what does it do?

SUM CLASSIFIER : $\text{Sign}(\vec{n} \cdot \vec{x} + b)$

HERE : $\text{SIGN}(\underbrace{\langle \phi(x) | W^\dagger(\theta) + W(\theta) | \phi(x) \rangle}_{\text{inner product?}} + b)$

$$\langle \phi(x) | W^\dagger + W | \phi(x) \rangle = \text{Tr} \left[\underbrace{W^\dagger + W}_A \overbrace{|\phi(x)\rangle\langle\phi(x)|}_B \right] = (A, B)_{\text{Fr}}$$

Det. $(\vec{w})_\alpha = \text{Tr}[W^\dagger + W P_\alpha]$; P_α - Pauli string $\alpha \in [0..4^{n-1}]$

$$(\vec{\phi}(\vec{x}))_\alpha = \text{Tr}(|\phi(x)\rangle\langle\phi(x)| P_\alpha)$$

$$\text{out}(\vec{x}) \hat{=} \text{SIGN}(\vec{w} \cdot \vec{\phi}(\vec{x}) + b)$$

the feature space is that of density operators...

What does it *do*?

$$[\vec{w}(\theta)]_{\alpha} = \text{tr} [W^{\dagger}(\vec{\theta}) \mathbf{f} W(\vec{\theta}) P_{\alpha}]$$

$$[\vec{\phi}(\vec{x})]_{\alpha} = \langle \Phi(\vec{x}) | P_{\alpha} | \Phi(\vec{x}) \rangle$$

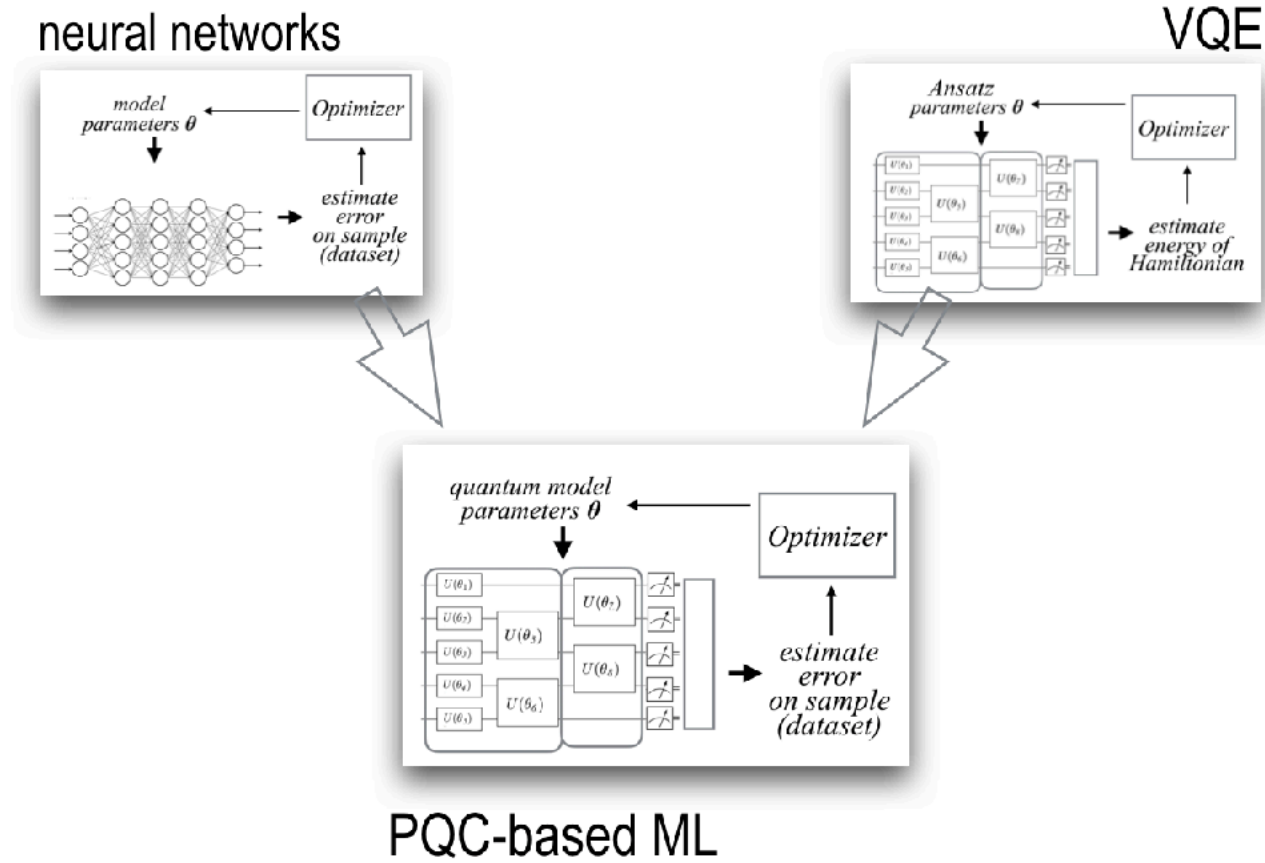
$$\text{out}(\vec{x}) \approx \text{sign}(\vec{w} \cdot \vec{\Phi}(\vec{x}) + b)$$

-limitations on the model
come into play here...
-not *all hyperplanes*
reachable...

-not maximal margin
attained!

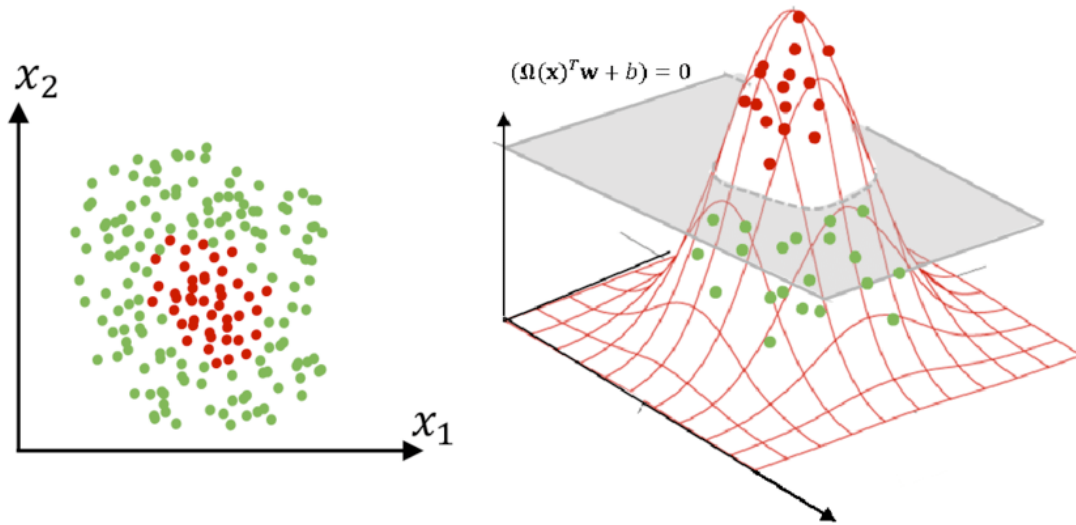
BECAUSE $W(\theta)$ & f ARE
RESTRICTED.

Note the explicit model is much like training a NN/VQE





-but with a connection with a well-understood classical model

Illustration of quantum decision boundaries



Two slices of quantum kernels:



Quantum advantage, and advantage for (near term) quantum

-for quantum advantage: *useful* and *classically* hard

-for advantage for near-term quantum: *useful* and *doable*

Quantum advantage, and advantage for (near term) quantum

-for quantum advantage: *useful* and *classically* hard

-useful: remains to be seen;

- almost all models useful in some settings; here when data has complex correlations.
- Bleeding edge reasearch:
- theory for ML is difficult;
- QCs just becoming large enough for experiments

Quantum advantage, and advantage for (near term) quantum



-for quantum advantage: *useful* and *classically* hard

-classically hard:

- trivially there exist “BQP-hard” kernels (for deep circuits)
- for “functional problems” no hard separation results but; very likely hard.
- more interestingly; likely hard in shallow circuit regime

Quantum advantage, and advantage for (near term) quantum

-for advantage for near-term quantum: *useful* and *doable*

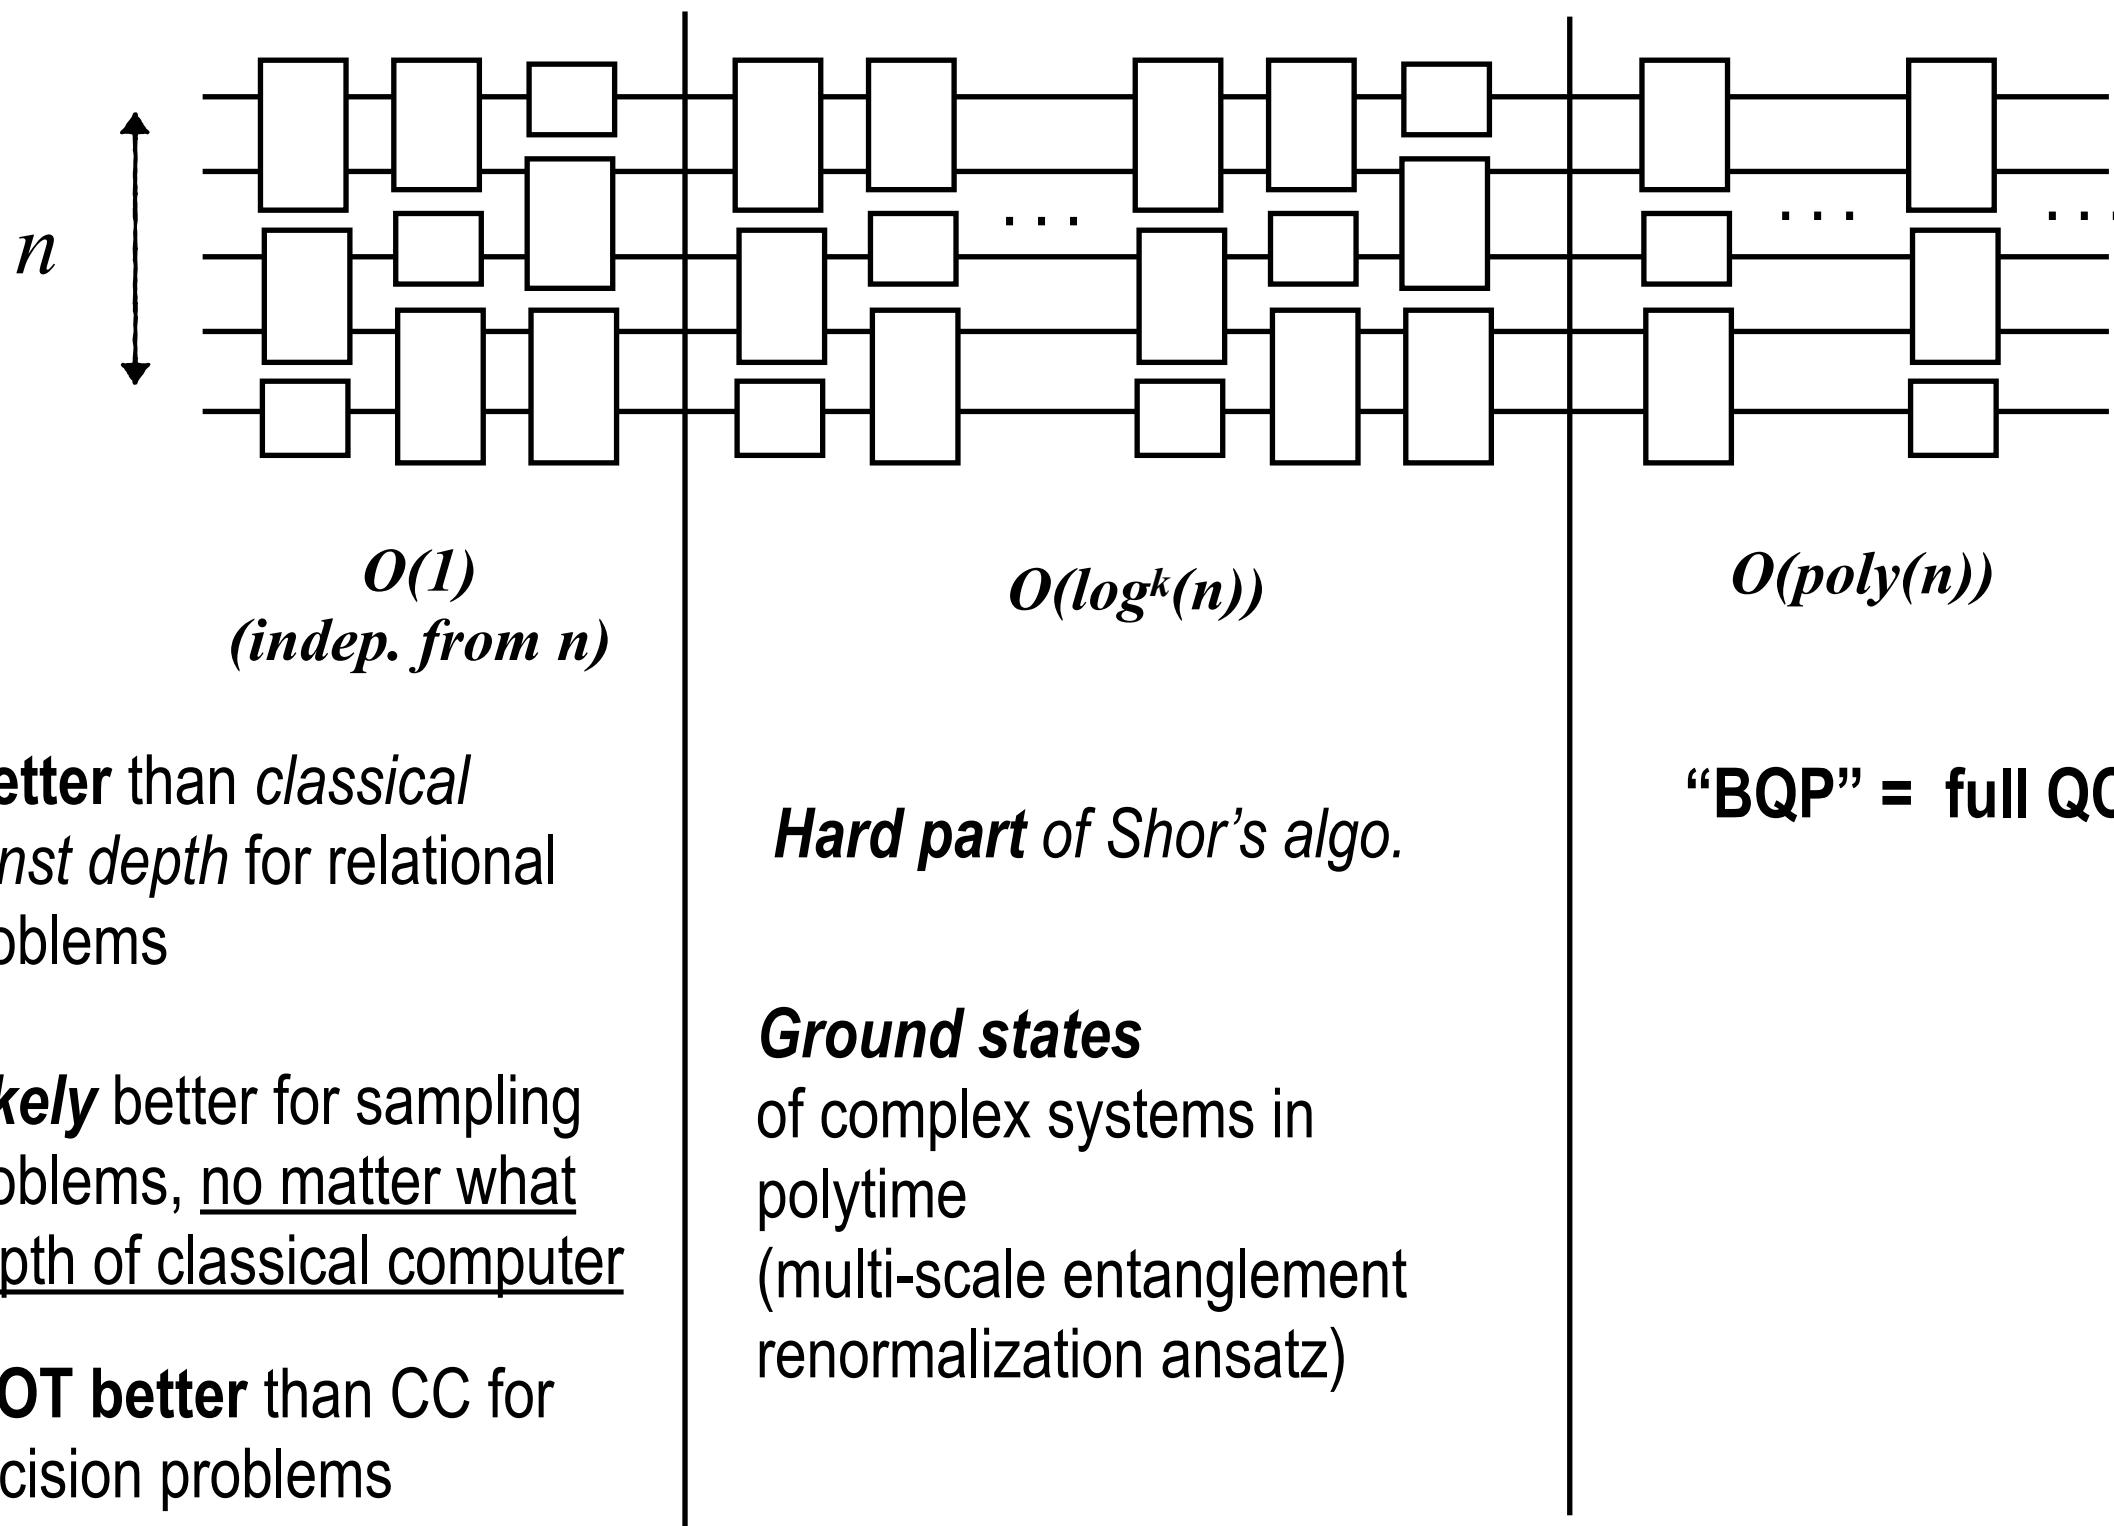


doable: makes sense with: ~100 qubits, limited depth, errors


Quantum advantage, and advantage for (near term) quantum

- 1) *~100 qubits - probably yes 2^{100} is interesting*
- 2) *depth?*
- 3) *noise?*

Recall Quantum depth complexity




can we have limited depth **and** classically hard?



L_m : full exact simulation of output of **log-many qubits** in constant depth, can be done in poly-time

This is the situation in chemistry with log-local Hamiltonians - depth must be at least log.



Not log-many!
No known efficient classical algorithm

- 1) *~100 qubits* ✓
- 2) *depth* ✓
- 3) *noise?*

Reasons for optimism:

- a) ML as signal-from-noise + source shifting*
- b) stochastic hypothesis families and noisy data (distinct from mathematical optimization)*
- c) brains are noisy :)*