

On the Principles of Differentiable Quantum Programming Languages



Shaopeng Zhu



Shih-Han Hung



Shouvanik Chakrabarti



Xiaodi Wu

PLDI 2020 - Language Design I



A Tale of Two Emerging Programming Languages

Heard of *Quantum* Programming Languages?

$|0\rangle, |1\rangle, \rho$

$\frac{\partial}{\partial x} [[P(x)]]$

Heard of *Differentiable* Programming Languages?

This talk is about the **happy marriage** of both:

Differentiable Quantum Programming Languages

However, it is not just a brain teaser but with **strong practical motivation!**

With the establishment of **Quantum Supremacy**,

IBM will soon launch a 53-qubit quantum computer

Frederic Lardinols @frederic / 8:00 am EDT • September 18, 2019

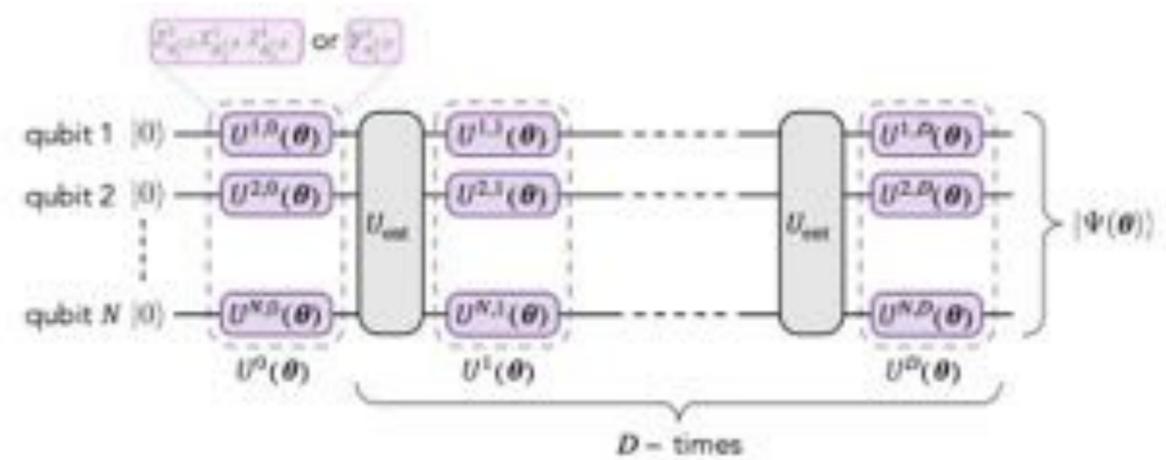


Google has reached quantum supremacy – here's what it should do next

TECHNOLOGY | ANALYSIS 26 September 2019
By Chelsea Whyte



Variational Quantum Circuits (VQC)



Its training requires **gradient-computation**
Infeasible for classical computation power

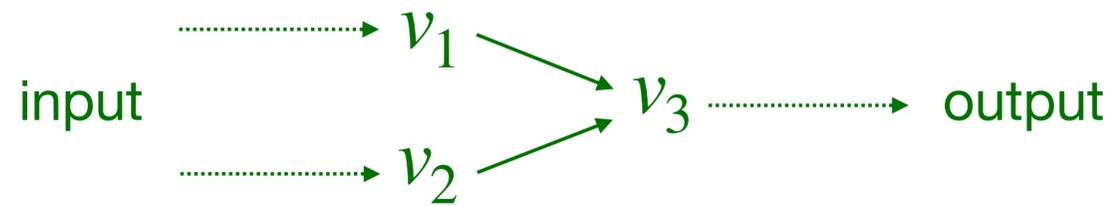
Thus,

using quantum programs to compute the gradients of quantum programs

is critical for the **scalability** of gradient-based quantum applications!

Classical vs Quantum 101

classical:



e.g.

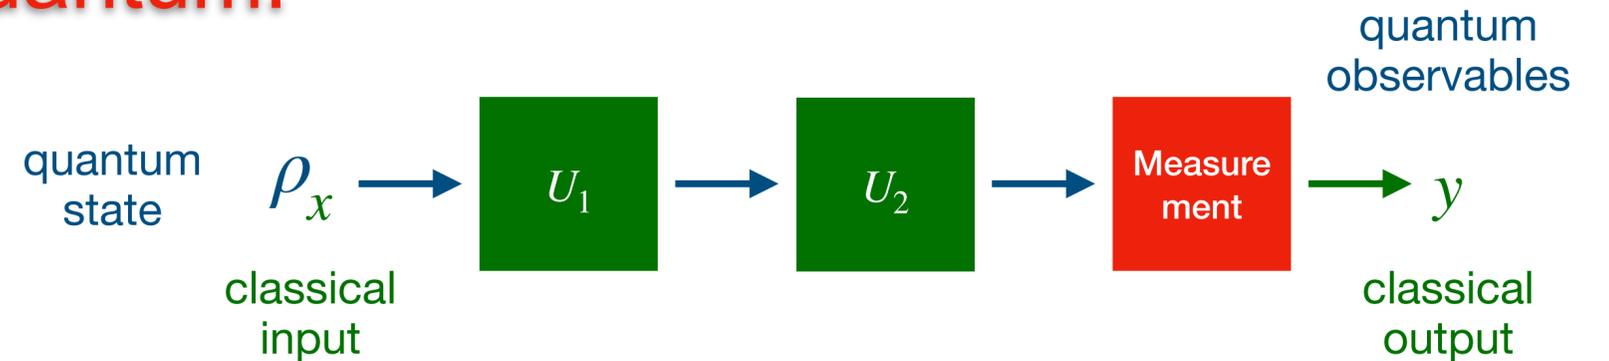
$$M \equiv v_3 = v_1 \times v_2$$
$$\frac{\partial}{\partial \theta} M \equiv v_3 = v_1 \times v_2;$$
$$\dot{v}_3 = \dot{v}_1 \times v_2 + v_1 \times \dot{v}_2$$

observations:

- actual state (v_i) = representation (v_i)
- all v_i are reals, thus *differentiable*
- store $v_1, v_2, v_3, \dot{v}_1, \dot{v}_2, \dot{v}_3$ at the same time

- chain-rule: $\frac{\partial v_3}{\partial \theta} = \frac{\partial v_3}{\partial v_1} \frac{\partial v_1}{\partial \theta} + \frac{\partial v_3}{\partial v_2} \frac{\partial v_2}{\partial \theta}$

quantum:

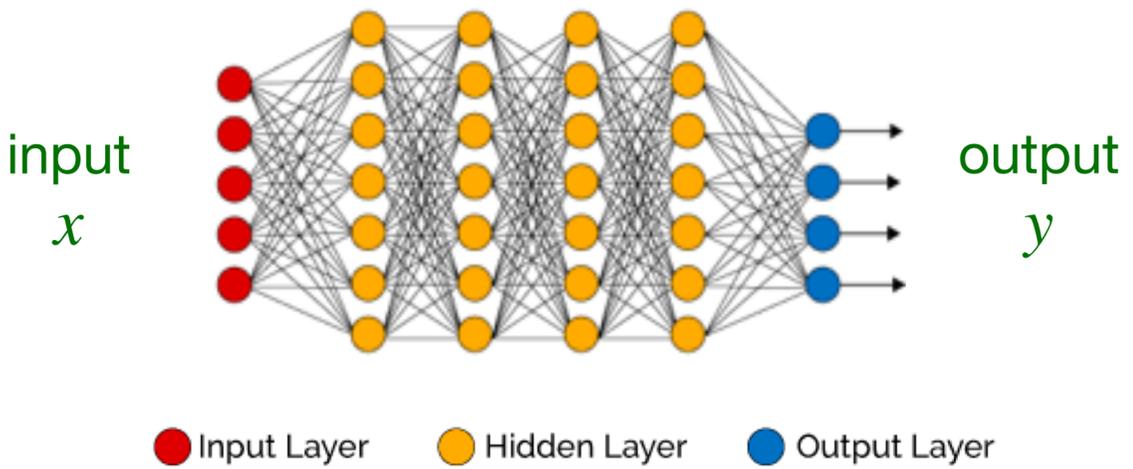


key differences:

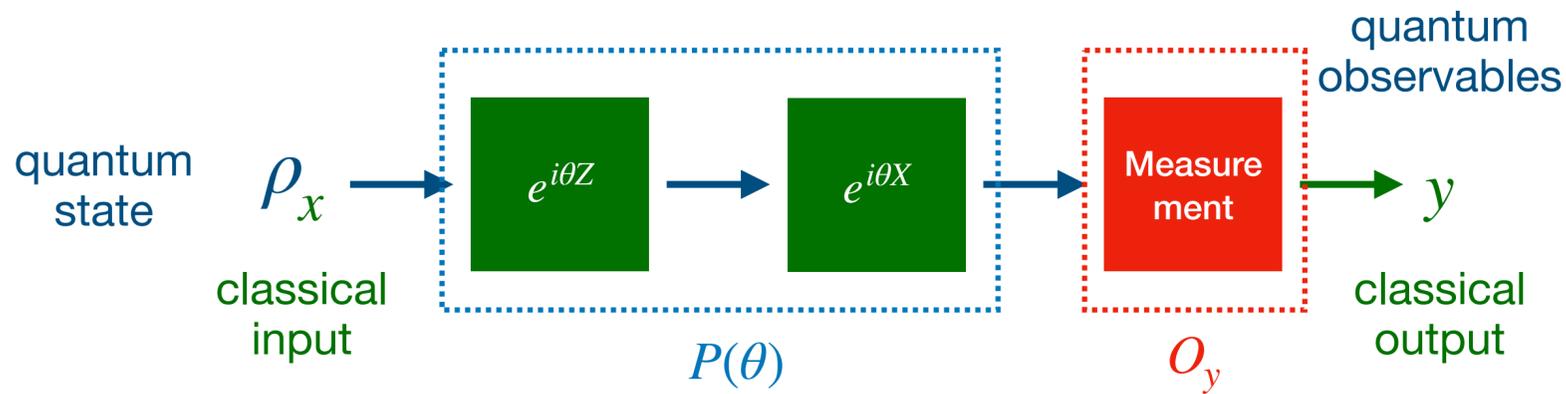
- actual state ρ is a **quantum** state; its **classical** representation is an exponential (in # of qubits) matrix.
- Unit operation U_1 takes **1 unit time** on **q.** machines; **classically** simulating $U_1 \rho U_1^\dagger$ takes **exponential time**.
- a priori unclear $\frac{\partial \rho}{\partial \theta}$ for both ρ and θ part
- cannot store all intermediate ρ due to **no-cloning**
- hard to make sense of chain-rules

Classical vs Quantum 101: cont'd

Classical Neural Networks (CNNs)

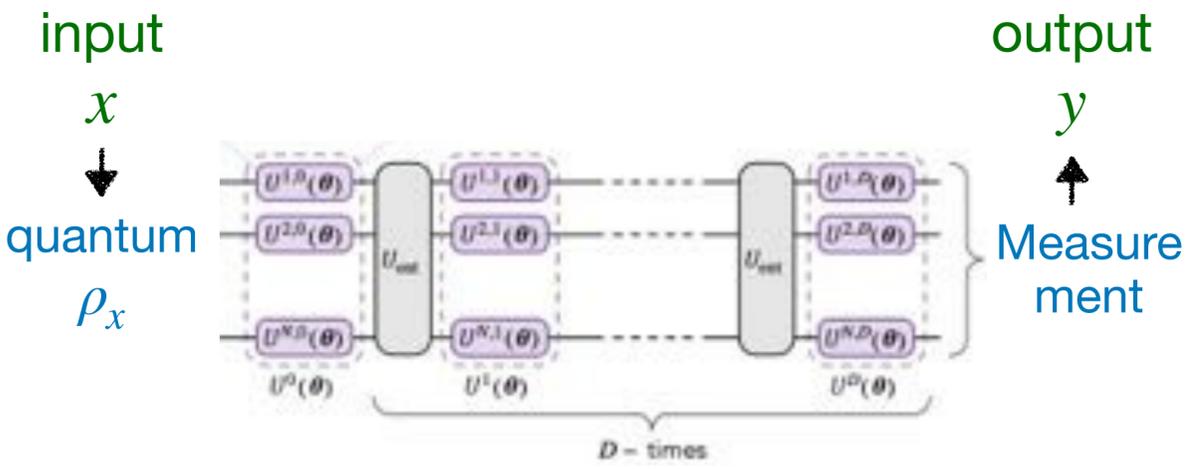


Variational Quantum Circuits (VQCs)



Replace $x \rightarrow y$ (classical) by

$x \rightarrow \rho_x \rightarrow y$ (quantum) w/ potential speedups



Variational Quantum Circuits (VQCs)

Training of VQCs, similar to CNNs, will optimize the **loss** functions

$$\text{loss} = L(x, y, \theta) \quad \text{for VQC } P(\theta) \text{ w/ d. semantics } [[P(\theta)]]$$

$$y(x, \theta) = \text{Tr}(O_y [[P(\theta)]](\rho_x))$$

quantum state ρ_x quantum observable O_y

Gradient $\frac{\partial L}{\partial \theta}$ can be computed from $\frac{\partial y}{\partial \theta}$ exponential classical cost

\exists q. gadget for simple P w/o formal formulation and program features

Conceptual challenges from the beginning and more to come!

Contributions

“Deep Learning est mort. Vive Differentiable Programming!”

----- Yann LeCun

- **Formal Formulation** of **Differentiable Quantum** Programming Languages:
 - *basic concepts*: **parameterized quantum programs, semantics of differentiation**
 - *code-transformation*: **two-stage code-transformation, a logic proving its correctness**
 - *features*: **support controls, compositions, also w/ resource efficiency**
- **Implementation** of a prototype in OCaml and benchmark tests on representative cases
 - *quantum neuro-symbolic application*: **parameterized quantum programs in machine learning**
 - *resource efficiency*: **empirically demonstrated efficiency for representative cases**



Differentiable Quantum Prog-Lang:
[github:/LibertasSpZ/adcompile](https://github.com/LibertasSpZ/adcompile)



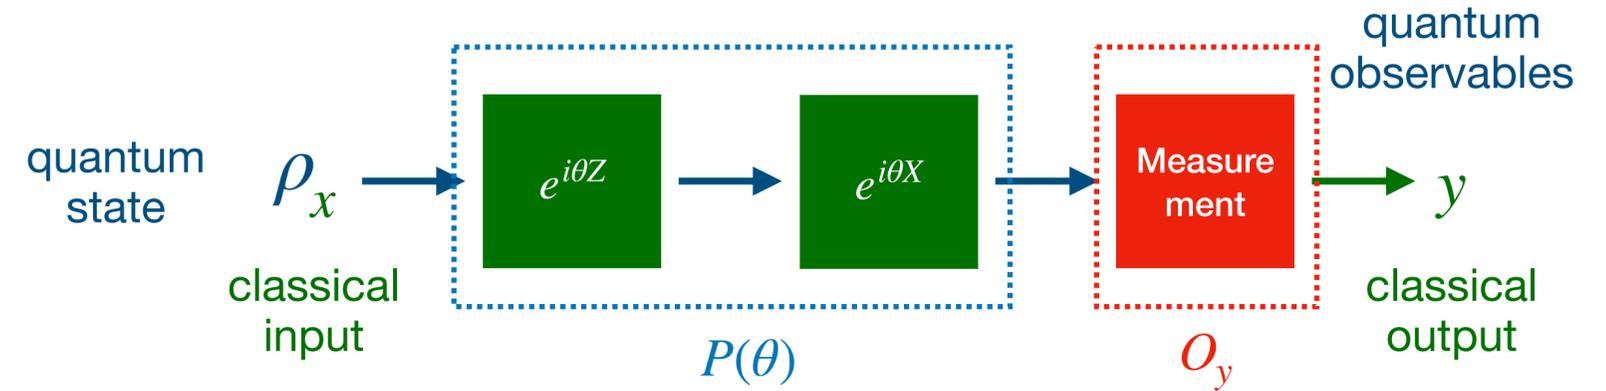
Formal Formulations:

Variational Quantum Circuits (VQCs)

- Use **parameterized quantum while-language** to formulate $P(\theta)$ with *classically parameterized* Pauli rotation gates: e.g.,

$$e^{i\theta Z}, e^{i\theta X \otimes X}$$

which forms a *universal* gate set and can be readily implemented on near-term quantum machines.



- Model **quantum observable** on $P(\theta)$ by **Observable Semantics**

$$\llbracket (O_y, \rho_x) \rightarrow P(\theta) \rrbracket \equiv \text{Tr}(O_y \llbracket P(\theta) \rrbracket (\rho_x)) \quad \text{matches exactly the observable quantum output}$$

Note that it will serve as (1) the target to differentiate; and (2) the read-out of any quantum programs.

- Make sense of $\frac{\partial}{\partial \theta} P(\theta)$ computing the derivative of $P(\theta)$ (**Differential Semantics**)

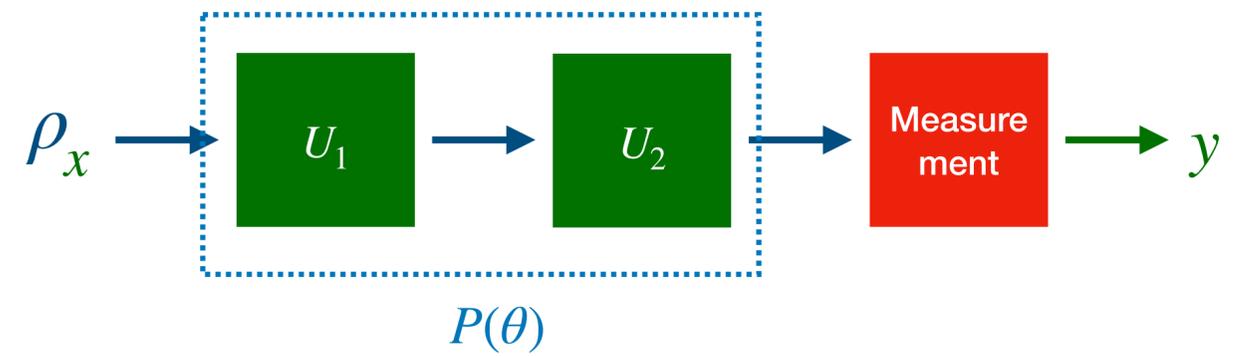
$$\llbracket (O_y, \rho_x) \rightarrow \frac{\partial}{\partial \theta} P(\theta) \rrbracket = \frac{\partial}{\partial \theta} \llbracket (O_y, \rho_x) \rightarrow P(\theta) \rrbracket$$

$$\text{one } \frac{\partial}{\partial \theta} P(\theta) \text{ for any } O_y \text{ and } \rho_x$$

strong requirement: achievable and critical

Formal Formulations: cont'd

Move to code-transformation and construction of $\frac{\partial}{\partial \theta} P(\theta)$ based on previous formulations



$$P(\theta) \equiv U_1(\theta); U_2(\theta);$$

$$\frac{\partial}{\partial \theta} M \equiv v_3 = v_1 \times v_2;$$

$$\dot{v}_3 = \dot{v}_1 \times v_2 + v_1 \times \dot{v}_2$$

classical analogue

$$\frac{\partial}{\partial \theta} P(\theta) \equiv \frac{\partial}{\partial \theta} U_1(\theta); U_2(\theta); + U_1(\theta); \frac{\partial}{\partial \theta} U_2(\theta); ??$$

Making sense of “+”:

classical “+”: run both $\dot{v}_1 \times v_2$ and $v_1 \times \dot{v}_2$ on the **input**, and then **sum** the outputs

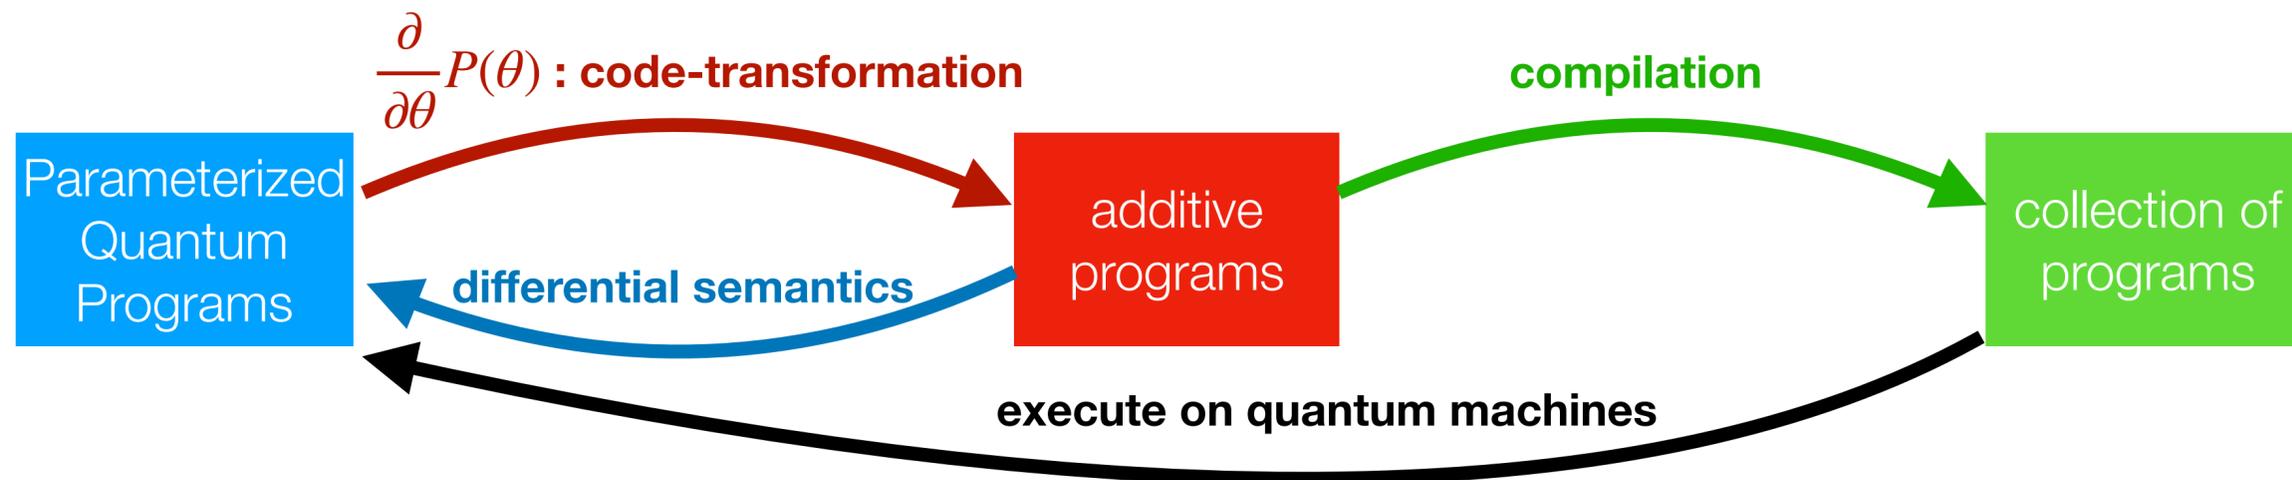
quantum “+”: hope to do the same. However, the **input cannot be cloned**, and will be **consumed** for each

Thus, $\frac{\partial}{\partial \theta} P(\theta)$ needs to be a **collection** of programs running on **copies** of the input state —> **complication!**

Ideally, (1) hope to have a **similar code-transformation like classical** for intuition and implementation, but still able to keep track of the right collection of programs in an efficient way.

(2) hope to control the **size** of the collection —> **# of copies** of the input state for efficiency.

Two-stage Code-Transformation



additive programs

Quantum While-Programs + Sum Operator

```

P(θ) ::= abort[q̄] | skip[q̄] | q := |0⟩ | q̄ := U(θ)[q̄] |
        P1(θ); P2(θ) | case M[q̄] = m → Pm(θ) end |
        while(T) M[q̄] = 1 do P1(θ) done | P1(θ) + P2(θ)
    
```

Operational Semantics of Sum

```

⟨P1(θ*) + P2(θ*), ρ⟩ → ⟨P1(θ*), ρ⟩,
⟨P1(θ*) + P2(θ*), ρ⟩ → ⟨P2(θ*), ρ⟩
    
```

exactly match the intuition of “+” for differential semantics

compilation

Output a collection of quantum programs while keeping the size small

$$\frac{\partial}{\partial \theta} P(\theta) \equiv \frac{\partial}{\partial \theta} U_1(\theta); U_2(\theta); + U_1(\theta); \frac{\partial}{\partial \theta} U_2(\theta);$$

$$\text{Compile}\left(\frac{\partial}{\partial \theta} P(\theta)\right) \equiv \{ \left[\frac{\partial}{\partial \theta} U_1(\theta); U_2(\theta);, U_1(\theta); \frac{\partial}{\partial \theta} U_2(\theta); \right], \text{size } 2 \}$$

general compilation follows the same intuition but more complicated :

```

(Atomic) Compile(P(θ)) ≡ {P(θ)},
        if P(θ) ≡ abort[v̄] | skip[v̄] | q := |0⟩
        | v̄ := U(θ)[v̄].
(Sequence) Compile(P1(θ); P2(θ)) ≡
    {
        {abort}, if Compile(P1(θ)) = {abort};
        {abort}, if Compile(P2(θ)) = {abort};
        {Q1(θ); Q2(θ) : Qb(θ) ∈ Compile(Pb(θ))},
        otherwise.
    }
(Case m) Compile(case) ≡ FB(case), described in Fig.3b.
(While(T)) Compile(while(T)) : use (Case m) and (Sequence).
(Sum) Compile(P1(θ) + P2(θ)) ≡
    {
        Compile(P1(θ)) || Compile(P2(θ)), if ∀ b ∈ {1, 2},
        Compile(Pb(θ)) ≠ {abort};
        Compile(P1(θ)), if Compile(P2(θ)) = {abort},
        Compile(P1(θ)) ≠ {abort};
        Compile(P2(θ)), if Compile(P1(θ)) = {abort},
        Compile(P2(θ)) ≠ {abort};
        {abort}, otherwise
    }
    
```

Code-Transformation

(Trivial)	$\frac{\partial}{\partial \theta}(\underline{\text{abort}}[\bar{v}]), \frac{\partial}{\partial \theta}(\underline{\text{skip}}[\bar{v}]), \frac{\partial}{\partial \theta}(q := 0\rangle) \equiv \underline{\text{abort}}[\bar{v} \cup \{A\}].$
(Trivial-U)	$\frac{\partial}{\partial \theta}(\bar{v} := U(\theta)[\bar{v}]) \equiv \underline{\text{abort}}[\bar{v} \cup \{A\}], \text{ if } \theta_j \notin \theta.$
(1-qb)	$\frac{\partial}{\partial \theta}(q_1 := R_\sigma(\theta)[q_1]) \equiv \underline{A, q_1 := R'_\sigma(\theta)[A, q_1]}.$
(2-qb)	$\frac{\partial}{\partial \theta}(q_1, q_2 := R_{\sigma \otimes \sigma}(\theta)[q_1, q_2]) \equiv \underline{A, q_1, q_2 := R'_{\sigma \otimes \sigma}(\theta)[A, q_1, q_2]}.$
(Sequence)	$\frac{\partial}{\partial \theta}(\underline{S_1(\theta); S_2(\theta)}) \equiv \underline{(S_1(\theta); \frac{\partial}{\partial \theta}(S_2(\theta)))} + \underline{(\frac{\partial}{\partial \theta}(S_1(\theta)); S_2(\theta))}.$
(Case)	$\frac{\partial}{\partial \theta}(\underline{\text{case } M[q] = m \rightarrow S_m(\theta) \text{ end}}) \equiv \underline{\text{case } M[\bar{q}] = m \rightarrow \frac{\partial}{\partial \theta}(S_m(\theta)) \text{ end}}.$
(while ^(T))	Use (Case) and (Sequence).
(S-C)	$\frac{\partial}{\partial \theta}(\underline{S_1(\theta)} + \underline{S_2(\theta)}) \equiv \underline{\frac{\partial}{\partial \theta}(S_1(\theta))} + \underline{\frac{\partial}{\partial \theta}(S_2(\theta))}.$

$\frac{\partial}{\partial \theta}P(\theta) : \text{code-transformation}$



We develop a sound logic to prove its correctness.

Similar to classical for the convenience of compiler implementation!

(1) modify existing *phase-shift* rule using two-circuit-difference to one circuit with super-posed control for composition and efficiency.

(2) the proof relies on the **strong** requirement in differential semantics

$$\begin{aligned} \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(\underline{S_0(\theta); S_1(\theta)}) \rrbracket &= \llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)); S_1(\theta) \rrbracket \\ &+ \llbracket (O, \rho) \rightarrow S_0(\theta); \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket. \end{aligned}$$

$$\llbracket (O, \rho) \rightarrow S_0(\theta); \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket = \llbracket (O, \llbracket S_0(\theta) \rrbracket(\rho)) \rightarrow \frac{\partial}{\partial \theta}(S_1(\theta)) \rrbracket$$

make use of the premises of $\frac{\partial}{\partial \theta}S_0(\theta)$,

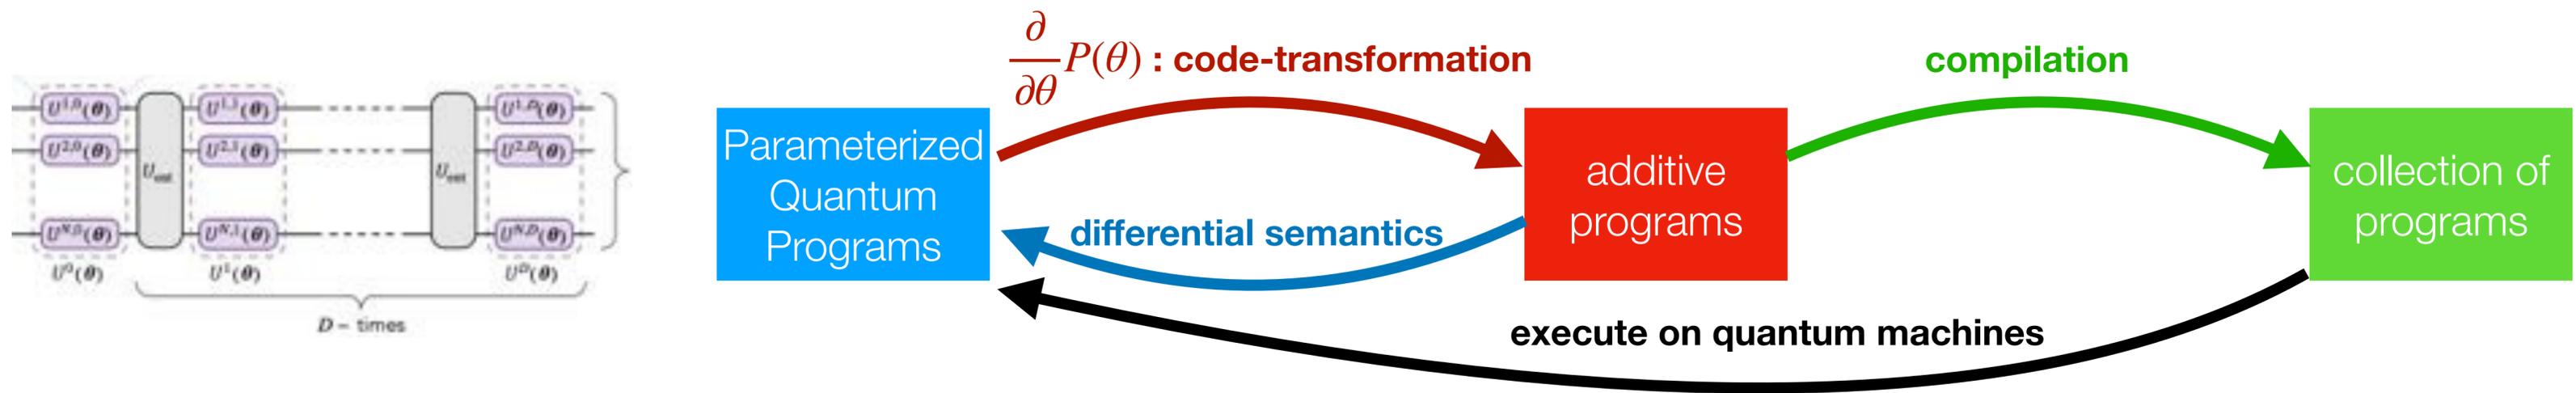
$$\llbracket (O, \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)); S_1(\theta) \rrbracket = \llbracket (\llbracket S_1(\theta) \rrbracket^*(O), \rho) \rightarrow \frac{\partial}{\partial \theta}(S_0(\theta)) \rrbracket$$

$\frac{\partial}{\partial \theta}S_1(\theta)$, for **different** (O, ρ) pairs

(3) support **case** unconditionally, better than the classical case

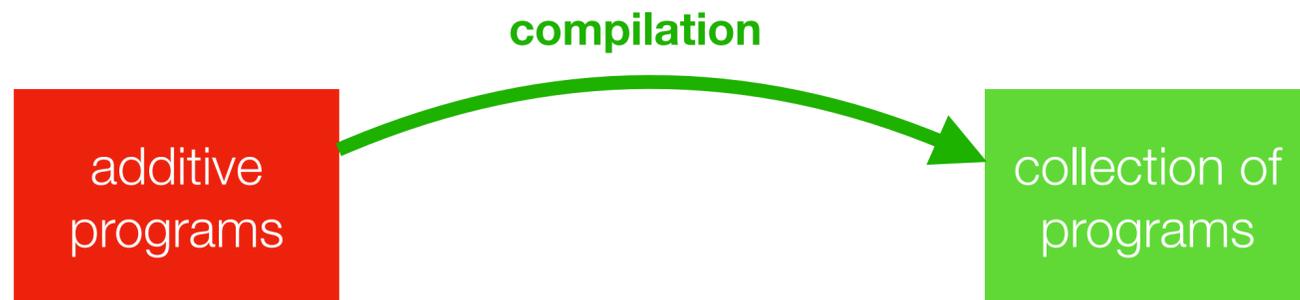
support **bounded loop**, matching the classical case [Plotkin, POPL'18]

Recap ...



- **Formal Formulation** of **Differentiable Quantum** Programming Languages:
 - *basic concepts*: parameterized quantum programs, semantics of differentiation
 - *code-transformation*: two-stage code-transformation, a logic proving its correctness
 - *features*: support controls, compositions, also w/ resource efficiency
- **Implementation** of a prototype in OCaml and benchmark tests on representative cases
 - *quantum neuro-symbolic application*: parameterized quantum programs in machine learning
 - *resource efficiency*: empirically demonstrated efficiency for representative cases

Resource Estimation



Generally **EXP** (in # of “+”s) programs after compilation

However, for relevant $\frac{\partial}{\partial \theta} P(\theta)$, we show the # is **bounded** by **occurrence count** of θ
 which basically counts the # of appearances of θ

$$P(\theta) \equiv U_1(\theta); U_2(\theta); \quad \text{occurrence count 2}$$

$$\text{Compile}\left(\frac{\partial}{\partial \theta} P(\theta)\right) \equiv \left\{ \left[\frac{\partial}{\partial \theta} U_1(\theta); U_2(\theta); \right], \left[U_1(\theta); \frac{\partial}{\partial \theta} U_2(\theta); \right] \right\}, \text{ size 2}$$

Why occurrence could be a reasonable quantity?

$$v_3 = v_1 \times v_2;$$

$$\dot{v}_3 = \dot{v}_1 \times v_2 + v_1 \times \dot{v}_2$$

occurrence count 2: $v_1(\theta), v_2(\theta)$ in v_3

instead of **extra initial states**, classically one needs 2 **extra registers** to store \dot{v}_1, \dot{v}_2

Definition 7.1. The “Occurrence Count for θ_j ” in $P(\theta)$, denoted $\text{OC}_j(P(\theta))$, is defined as follows:

1. If $P(\theta) \equiv \text{abort}[\bar{v}] | \text{skip}[\bar{v}] | q := |0\rangle$ ($q \in \bar{v}$), then $\text{OC}_j(P(\theta)) = 0$;
2. $P(\theta) \equiv U(\theta)$: if $U(\theta)$ trivially uses θ_j , then $\text{OC}_j(P(\theta)) = 0$; otherwise $\text{OC}_j(P(\theta)) = 1$.
3. If $P(\theta) \equiv U(\theta) = P_1(\theta); P_2(\theta)$ then $\text{OC}_j(P(\theta)) = \text{OC}_j(P_1(\theta)) + \text{OC}_j(P_2(\theta))$.
4. If $P(\theta) \equiv \text{case } M[\bar{q}] = m \rightarrow P_m(\theta) \text{ end}$ then $\text{OC}_j(P(\theta)) = \max_m \text{OC}_j(P_m(\theta))$.
5. If $P(\theta) \equiv \text{while}^{(T)} M[\bar{q}] = 1 \text{ do } P_1(\theta) \text{ done}$ then $\text{OC}_j(P(\theta)) = T \cdot \text{OC}_j(P_1(\theta))$.

$P(\theta)$	$\text{OC}(\cdot)$	$ \# \frac{\partial}{\partial \theta}(\cdot) $	#gates	#lines	#layers	#qb
$\text{QNN}_{M,i}$	24	24	165	189	3	18
$\text{QNN}_{M,w}$	56	24	231	121	5	18
$\text{QNN}_{L,i}$	48	48	363	414	6	36
$\text{QNN}_{L,w}$	504	48	2079	244	33	36
$\text{VQE}_{M,i}$	15	15	224	241	3	12
$\text{VQE}_{M,w}$	35	15	224	112	5	12
$\text{VQE}_{L,i}$	40	40	576	628	5	40
$\text{VQE}_{L,w}$	248	40	1984	368	17	40
$\text{QAOA}_{M,i}$	18	18	120	142	3	18
$\text{QAOA}_{M,w}$	42	18	168	94	5	18
$\text{QAOA}_{L,i}$	36	36	264	315	6	36
$\text{QAOA}_{L,w}$	378	36	1512	190	33	36

Benchmark tests on **typical** instances

Quantum Neuro-Symbolic Application

“Deep Learning est mort. Vive Differentiable Programming!”

----- Yann LeCun

$$Q(\Gamma) \equiv R_X(\gamma_1)[q_1]; R_X(\gamma_2)[q_2]; R_X(\gamma_3)[q_3]; R_X(\gamma_4)[q_4]; \\ R_Y(\gamma_5)[q_1]; R_Y(\gamma_6)[q_2]; R_Y(\gamma_7)[q_3]; R_Y(\gamma_8)[q_4]; \\ R_Z(\gamma_9)[q_1]; R_Z(\gamma_{10})[q_2]; R_Z(\gamma_{11})[q_3]; R_Z(\gamma_{12})[q_4],$$

(no control) $P_1(\Theta, \Phi) \equiv Q(\Theta); Q(\Phi)$.

(w/ control) $P_2(\Theta, \Phi, \Psi) \equiv Q(\Theta); \text{case } M[q_1] = 0 \rightarrow Q(\Phi) \\ 1 \rightarrow Q(\Psi)$.

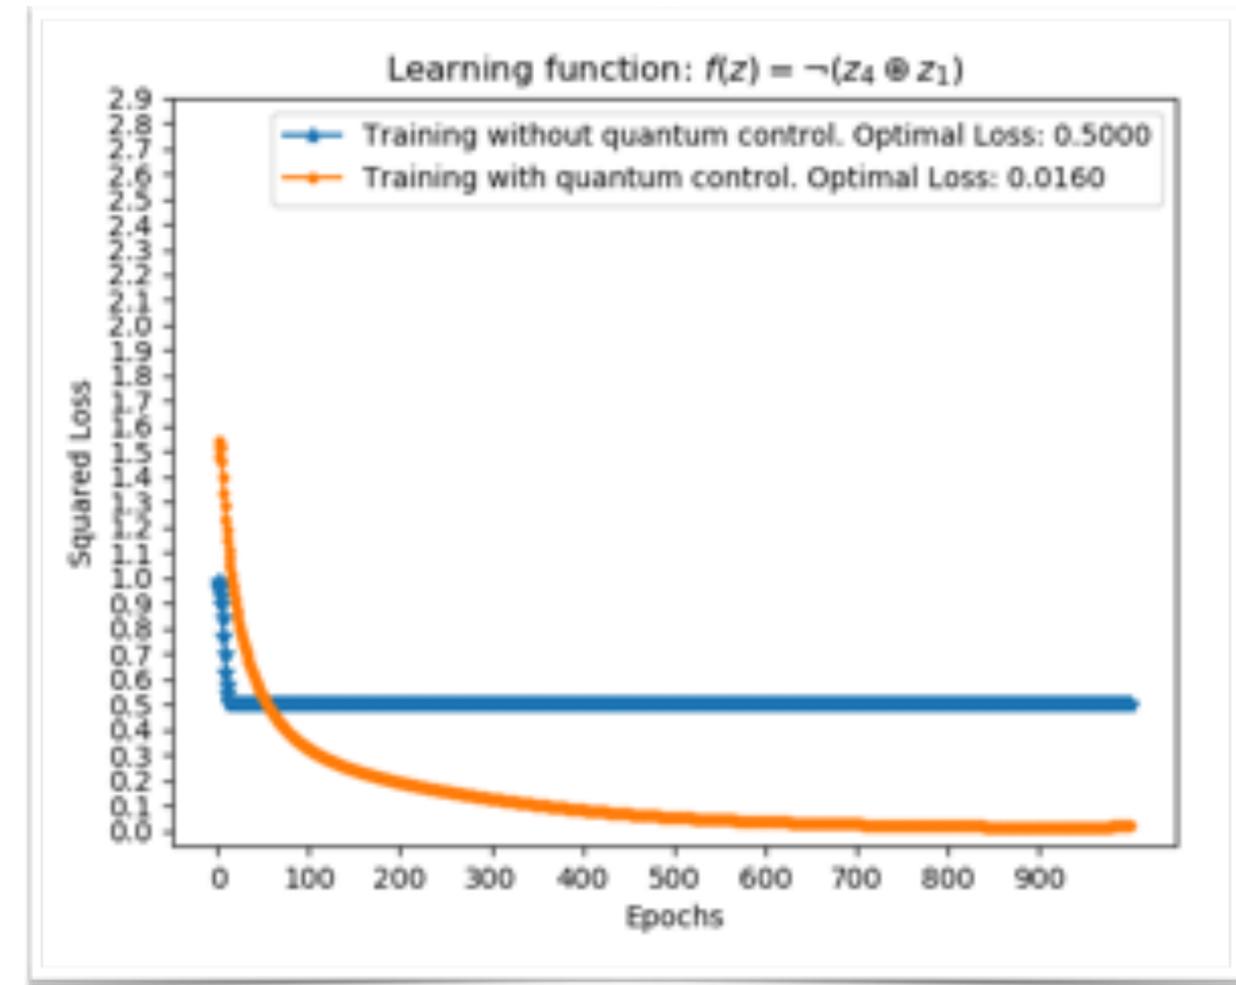
Note that $P_1(\Theta, \Phi)$ and $P_2(\Theta, \Phi, \Psi)$ run the same # of gates.

Simple **classification** task w/ ground truth

$$f(z) = \neg(z_1 \oplus z_4), z = z_1 z_2 z_3 z_4 \in \{0,1\}^4$$

via the following **square loss** function

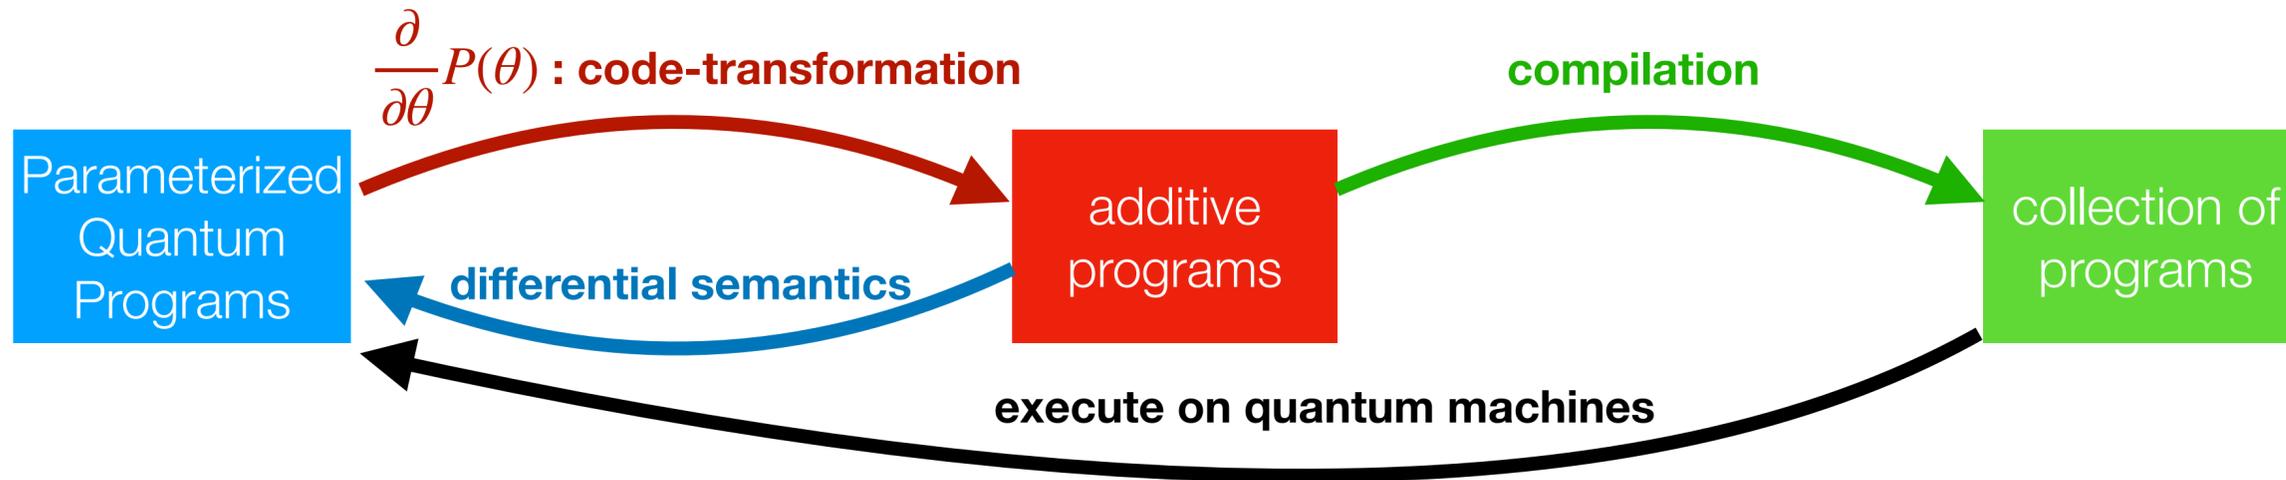
$$\text{loss} = \sum_{z \in \{0,1\}^4} 0.5 * (l_{\theta}(z) - f(z))^2$$



Training of P_2 is conducted on our implementation, whereas the training of P_1 is conducted on prior art (e.g., PennyLane).

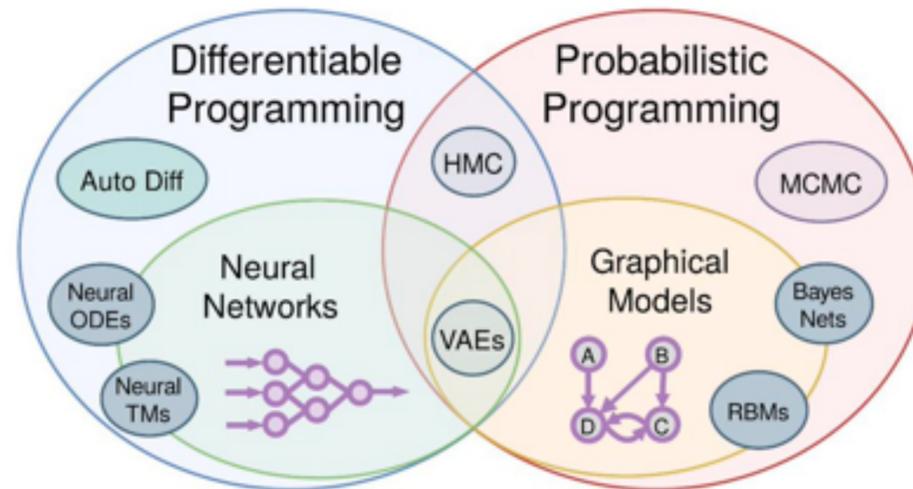
The use of **quantum control** could be **significant** for near-term.

Conclusions



open questions:
functional PL,
q. neuro-symbolic apps

even if you are
only interested
in classical



As a generalization of both,
our findings might provide hints of **unifying**
differentiable and **probabilistic** programming!



THANK YOU!

Differentiable Quantum Prog-Lang:
[github:/LibertasSpZ/adcompile](https://github.com/LibertasSpZ/adcompile)

