

# Optimization for Life

Fei Sun

November 2021

# Purpose of the talk

- I'm looking for collaborators.
- You don't need to agree with my philosophy.
- If you feels that some of the proposed directions are of interest to you, let's talk.

# Disclaimer

- This is NOT a scientific presentation
- This presentation contains many
  - Hypotheses
  - Projections
  - Extrapolations

# Outline

- What life do we optimize?
- Immortality is not a dream
- The next AI revolution
- A practical way to become the next NVIDIA
- Research directions

What Life Do We Optimize ?

# Observation

- I can **control** myself only
  - I **decide** what to do for myself, though with **external constraints**
    - What to have for lunch?
    - What career to take?
    - Which company to work for?
- I can only **inference** others
  - If I hope others to do something, others have to agree for the benefit of themselves
  - Why do I want to inference others?
    - Optimize external constraints so I can achieve my decision better
- Then who is the most important person in the world?

# Observation

- I can **control** myself only
  - I **decide** what to do for myself, though with **external constraints**
    - What to have for lunch?
    - What career to take?
    - Which company to work for?
- I can only **inference** others
  - If I hope others to do something, others have to agree for the benefit of themselves
  - Why do I want to inference others?
    - Optimize external constraints so I can achieve my decision better
- Then who is the most important person in the world?
  - **Me**

# Optimize for Current Life

- Maximize your **whole-life** reward with your whole-life effort\*:

- $R = \max(\int_t^T f(t) dt)$  (1)

- $E = \int_t^T \hat{f}(t) dt$  (2)

- $f(t') = \int_0^{t'} F(\hat{f}(t), t') dt \quad (t' < T)$

- *t*: current time
- *T*: current life
- *f*(*t*): immediate reward
- *f*<sup>^</sup>(*t*): immediate effort
- Optimize **self**, e.g.
  - Pursue financial freedom (remove financial external constraint)

\* The discussion assumes that humans are rational. It also contains reward hypothesis.



# Optimize for Current Life and Legacy

- Spare some utilities to after-life legacy:

- $R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt)$  (3)

- $E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt$  (4)

- $f(t') = \int_0^t F(\hat{f}(t), t') dt$  ( $t' < T$ )

- $g(t') = \int_0^T G(\hat{g}(t))dt$  ( $T < t'$ )

- $t$ : *current time*
- $T$ : *current life*
- $f(t)$ : *immediate reward*
- $\hat{f}(t)$ : *immediate effort*
- $g(t)$ : *immediate legacy reward*
- $\hat{g}(t)$ : *immediate legacy effort*
- $\int_t^T g(t)dt$  is merged into  $\int_t^T f(t) dt$

# Optimize for Current Life and Legacy

- Spare some utilities to after-life legacy:

- $R = \max(\int_t^T f(t) dt + \int_T^\infty g(t) dt)$  (3)

- $E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t) dt$  (4)

- Define:  $\alpha = \int_t^T \hat{f}(t) dt$ ,  $\beta = \int_t^T \hat{g}(t) dt$

- Free-will choices of  $\alpha$  and  $\beta$ , e.g.

- $\alpha > 0, \beta > 0$ : care for offspring – descendants are your legacy
  - $\alpha > 0, \beta > 0$ : build long lasting artifacts – future generations are your legacy
    - Founders of a nation, inventors, scientists
  - $\alpha > 0, \beta = 0$ : degenerate to (2)
  - $\alpha = 0, \beta \gg 0$ : martyr – sacrifice for the better good of other people
  - $\alpha > 0, \beta < 0$ : traitor – willingly or unwillingly neglects legacy

- We no longer exist, why do we care about legacy?

- There is no feedback to benefit **self**
  - Only **belief** and humanity for a greater population
    - But what is their relations to **self**?

# Optimize for Current Life, Legacy, and After-Life

- Spare some utilities to after-life legacy:

- $R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt)$  (5)

- $E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt$  (6)

- $f(t') = \int_0^t F(\hat{f}(t), t') dt \quad (t' < T)$

- $g(t') = \int_0^T G(\hat{g}(t))dt \quad (T < t')$

- $h(t') = \int_0^T H(\hat{h}(t))dt \quad (T < t')$

- $t$ : current time
- $T$ : current life
- $f(t)$ : immediate reward
- $\hat{f}(t)$ : immediate effort
- $g(t)$ : immediate legacy reward
- $\hat{g}(t)$ : immediate legacy effort
- $h(t)$ : immediate after – life reward
- $\hat{h}(t)$ : immediate after – life effort
- $\int_t^T g(t)dt$  is merged into  $\int_t^T f(t) dt$

# Optimize for Current Life, Legacy, and After-Life

- Spare some utilities to after-life:

$$\bullet R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt) \quad (5)$$

$$\bullet E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt \quad (6)$$

- Define:  $\alpha = \int_t^T \hat{f}(t) dt$ ,  $\beta = \int_t^T \hat{g}(t)dt$ ,  $\gamma = \int_t^T \hat{h}(t)dt$

- Free-will choices of  $\alpha$ ,  $\beta$  and  $\gamma$ , e.g.

- $\gamma = 0$ : atheist – there is no after-life
- $\gamma > 0$ : Christian, Buddhist, etc. – strong **belief** that after-life exist
  - $\alpha, \beta, \gamma > 0$ : balanced perform actions for the purpose of better after-life
  - $\alpha, \beta \rightarrow 0, \gamma \gg 0$ : ascetic – sacrifice current life for better after-life
  - $\alpha = 0, \gamma \gg 0$ : martyr (religious), e.g. fundamentalist.

$$\bullet \int_t^T \hat{f}(t) dt \Rightarrow \int_t^T f(t) dt, \int_t^T \hat{g}(t)dt \Rightarrow \int_T^\infty g(t)dt$$

- They can be examined

$$\bullet \int_t^T \hat{h}(t)dt \Rightarrow ? \int_T^\infty h(t)dt$$

- It cannot be examined. It is only supported by **belief**. i.e.  $H$  may be zero.
- If  $H > 0$ , the effort still benefits self, though not in the current life

# Do You Spare More Effort On $\alpha$ , $\beta$ , Or $\gamma$ ?

- $R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt)$  (5)

- $E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt$  (6)

- Define:  $\alpha = \int_t^T \hat{f}(t) dt$ ,  $\beta = \int_t^T \hat{g}(t)dt$ ,  $\gamma = \int_t^T \hat{h}(t)dt$

- Let's pause a moment...

# Do You Spare More Effort On $\alpha$ , $\beta$ , Or $\gamma$ ?

- $R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt)$  (5)

- $E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt$  (6)

- Define:  $\alpha = \int_t^T \hat{f}(t) dt$ ,  $\beta = \int_t^T \hat{g}(t)dt$ ,  $\gamma = \int_t^T \hat{h}(t)dt$

- Everyone in history must make a choice for  $\alpha$ ,  $\beta$ , and  $\gamma$ , consciously or unconsciously ...

# Do We Spare More Effort On $\alpha$ , $\beta$ , Or $\gamma$ ?

$$\bullet R = \max(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt) \quad (5)$$

$$\bullet E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt \quad (6)$$

$$\bullet \text{Define: } \alpha = \int_t^T \hat{f}(t) dt, \beta = \int_t^T \hat{g}(t)dt, \gamma = \int_t^T \hat{h}(t)dt$$

• Everyone in history must make a choice for  $\alpha$ ,  $\beta$ , and  $\gamma$ , consciously or unconsciously ...

• Until now...

# It Is Time to Change The Equation

- $R = \infty$  (5)

- $E = \lim_{T \rightarrow \infty} \int_t^T \hat{f}(t) dt$  (6)

- Why now?
- How?



Immortality is not a dream

Let's walk through my thinking process

# The Seven Stages of Evolution\*

S0: Behavior indifferent to environment, evolution indifferent to environment

S1: Behavior indifferent to environment, evolution reactive to environment -> virus, some plants

S2: Behavior reactive to environment, evolution reactive to environment -> animals

S3: Behavior proactive to environment, evolution reactive to environment -> humans

**S4: Behavior proactive to environment, evolution proactive to environment**

S5: Behavior proactive to environment, evolution is morphed to self-evolution

S6: Behavior controls the environment, life is eternal

S7: Behavior is the environment

- S4 may take a long time... there is no chance for me to see later stages.... but something changed...

\* Some scribble of things

<https://feisun.org/2017/12/24/some-scribble-of-things/>

Let's ask ourselves this question:

Why we have not found  
extraterrestrial life?

# Fermi Paradox\*

- There are billions of stars in the Milky Way similar to the Sun.<sup>[3][4]</sup>
- With high probability, some of these stars have Earth-like planets in a circumstellar habitable zone.<sup>[5]</sup>
- Many of these stars, and hence their planets, are much older than the Sun.<sup>[6][7]</sup> If the Earth is typical, some may have developed intelligent life long ago.
- Some of these civilizations may have developed interstellar travel, a step humans are investigating now.
- Even at the slow pace of currently envisioned interstellar travel, the Milky Way galaxy could be completely traversed in a few million years.<sup>[8]</sup>
- And since many of the stars similar to the Sun are billions of years older, Earth should have already been visited by extraterrestrial civilizations, or at least their probes.<sup>[9]</sup>
- **However, there is no convincing evidence that this has happened.**<sup>[8]</sup>

\* [https://en.wikipedia.org/wiki/Fermi\\_paradox](https://en.wikipedia.org/wiki/Fermi_paradox)

# Von Neumann Probe / Bracewell Probe

- Self-replicating probes could exhaustively explore a galaxy the size of the Milky Way in as little as a million years.\*
- Why we haven't seen them?

\* [https://en.wikipedia.org/wiki/Fermi\\_paradox](https://en.wikipedia.org/wiki/Fermi_paradox)

# Drake Equation\*

$$N = R_* * f_p * n_e * f_l * f_i * f_c * L$$

where:

$N$  = the number of [civilizations](#) in our galaxy with which communication might be possible (i.e. which are on our current past [light cone](#));  
and

$R_*$  = the average rate of [star formation](#) in [our Galaxy](#)

$f_p$  = the fraction of those stars that have [planets](#)

$n_e$  = the average number of planets that can potentially support [life](#) per star that has planets

$f_l$  = the fraction of planets that could support life that actually develop life at some point

$f_i$  = the fraction of planets with life that actually go on to develop [intelligent](#) life (civilizations)

$f_c$  = the fraction of civilizations that develop a technology that releases detectable signs of their existence into space

$L$  = the length of time for which such civilizations release detectable signals into space<sup>[5]</sup>

\* [https://en.wikipedia.org/wiki/Drake\\_equation](https://en.wikipedia.org/wiki/Drake_equation)


# “Where Is Everybody?”

- Earth hosts the only life in the universe.
  - We'd better all be Christians
- Dark forest theory
  - People develop such gene based on evidence.
  - But why we haven't seen any evidence of interstellar battles?
- Dark domain (黑域)
  - Do we have evidence of dark domain?
  - It is stage 5, not stage 4



# The Universal Model of Evolution\*

- Stage 4 is the stage to explore the universe
  - Be forced to explore
- But we have not seen any, why?



\* Some scribble of things

<https://feisun.org/2017/12/24/some-scribble-of-things/>

# Because: Stage 4 may be short

- Path 1: Most civilizations end at stage 4
  - Stage 4 is the **MOST** dangerous stage
  - World is ended before exploring the universe
- Path 2: People evolve to stage 5 very fast
  - Technology advances exponentially
  - Once we reach stage 5, we are immortal...
  - Once we reach stage 5, we don't want to be detected...
- **We only explore the universe in stage 4, but it is short**
  - **That's the reason that we do not see any extraterrestrial life**

# 永生的阶梯 - 刘慈欣

如果我说，**有史以来的所有人基本上都是平等的**——是有史以来，不是法国大革命以后。你有什么感觉？如果大部分人觉得荒谬，那是因为他们还没有见过更大的不平等，或者说，还没有出现这样不平等的技术条件。

你在人生的平原上走着走着，迎面遇到一堵墙，这墙向上无限高，向下无限深，向左边和右边都无限长。这墙是什么应该不难猜到。在过去的时代，平民可能走三四十年就遇到这堵墙，帝王和贵族可能走出七八十年才遇到，但他们之间相差一般不会超过五十年，如后面所述，这个差别微不足道。所有人在相差不到一个数量级的时间里遇到这堵墙，这是最大的平等，这堵墙就是上帝或大自然为人类社会设置的平等的底线。

**但随着技术的发展，有些人前面的死亡之墙要被拆掉了，人生的平原对于他们将无限广阔。**

But this is just a science fiction... is it so?

# Immortality May Be Possible

- We may be possible to see immortality in our lifetime!
  - It may suffer from the “overconfidence effect”.\*
- Will stage 4 be so short that it does not exist?
  - I don’t know
  - If it takes a long time, it doesn’t matter to **ME**.
- Where should I spend my effort?

$$R = \max\left(\int_t^T f(t) dt + \int_T^\infty g(t)dt + \int_T^\infty h(t)dt\right)$$

$$E = \int_t^T \hat{f}(t) dt + \int_t^T \hat{g}(t)dt + \int_t^T \hat{h}(t)dt$$

Or

$$R = \infty$$
$$E = \lim_{T \rightarrow \infty} \int_t^T \hat{f}(t) dt$$

\* [https://en.wikipedia.org/wiki/Overconfidence\\_effect](https://en.wikipedia.org/wiki/Overconfidence_effect)

# I am selfish

- If this is possible, we will see the greatest **inequality** in history.
- I don't want to be at the sideline.
- I want to make it possible. I want to speed it up.
- **And enjoy the benefit of creating it.**

# 永生阶梯第一级

- 永生阶梯的第一级就是活过五十年。
- Technology advances exponentially in the next 50 years so that the later stages become possible.
  - Direct technology is life science
- How can we help to speed up the process?
  - Start from what we can do

Note: I don't believe all stairs described in the novel

Let's design an AGI to work  
on the rest!

This is my real reason to work on AGI

# The Next AI Revolution




# Previous Works on AGI

- 自上而下思考实现AGI技术难点及可能方法
- 稀疏是通往AGI的必由之路
- AGI，从我做起
  
- This time we talk about something different

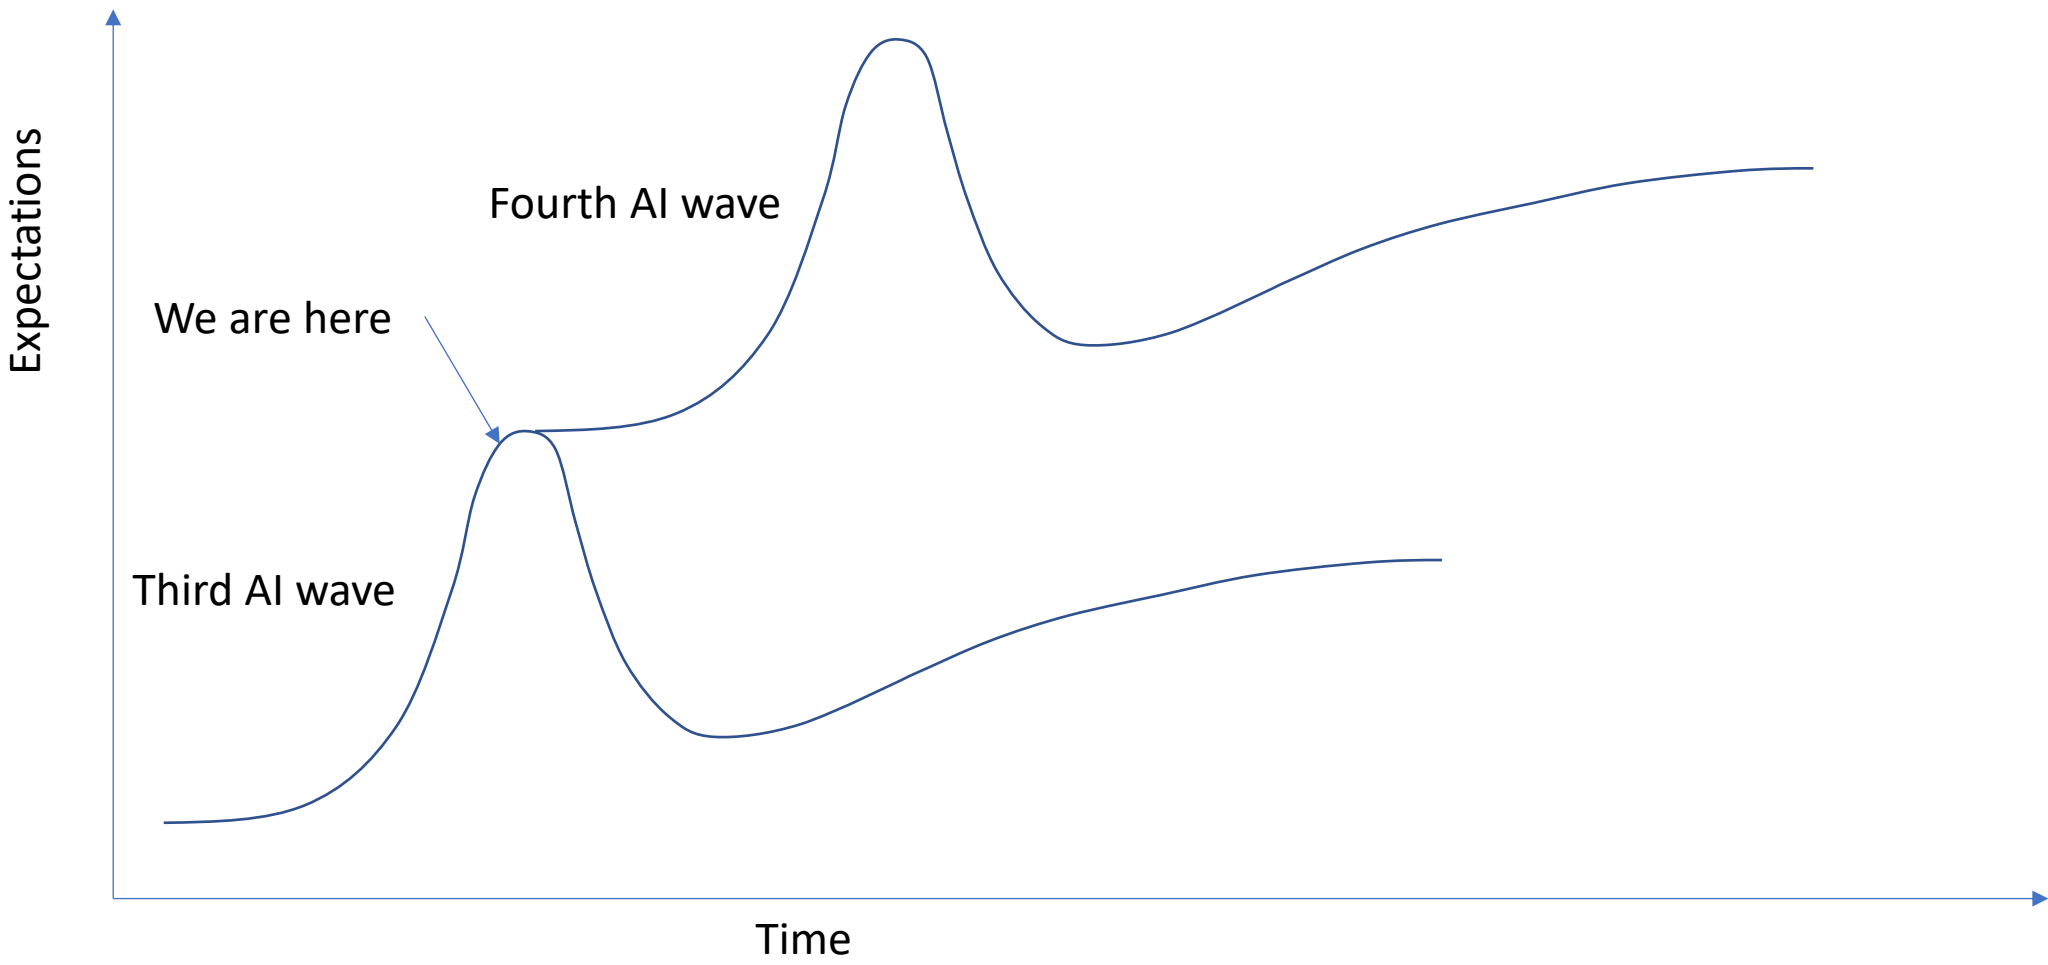
# Current Mainstream DL is Going Downhill

- Foundation model has swept through all DNN fields
  - Many cheer on the success
- However, it requires too much data, too much compute power
  - Its future potential is limited. It cannot lead us to AGI
  - Coz nobody can afford such expansion
- If we target AGI, we need to find another way.

# The Hype Cycle for the Third AI Wave



# The Hype Cycle for the Fourth AI Wave



# The Fourth AI wave

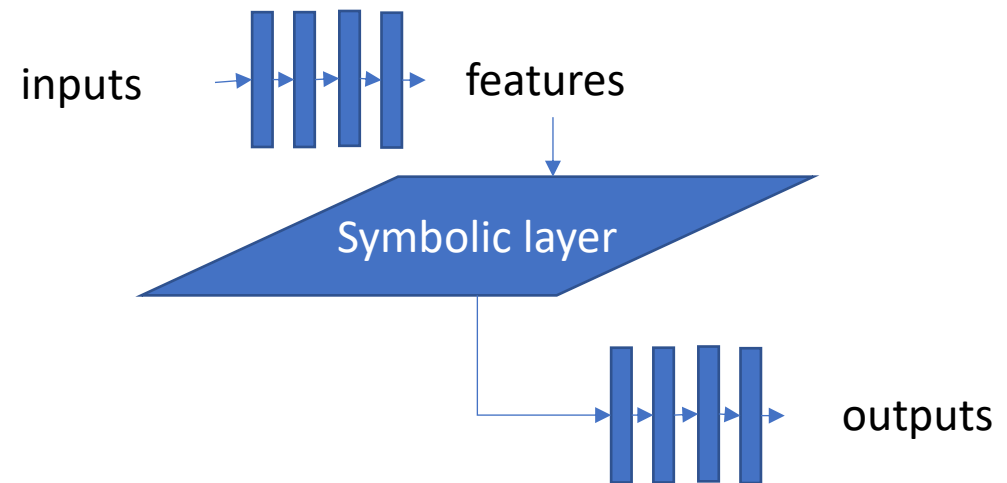
- Abstract symbolic layer – neurosymbolic computing
- Control flow model
- Revolutionized reinforcement learning
  - Apply DNN based control flow

# We Need to Raise the Abstraction Layer

- The existing DNNs mix low level features with the high level methods
  - Features are data dependent
  - Methods are data independent
- Insert an abstraction layer to decouple features and methods
- -> Insert a **symbolic** layer
- -> Neurosymbolic AI
  - Combine connectionist AI and symbolic AI
  - Getting more attention since 2020
- Composition is built-in

# One View of Neurosymbolic AI

- Connectionist AI extracts features and abstract to symbols
- Connectionist AI conclude methods from symbols



How to explore methods?

# Control Flow Model

We want to do this:

```
if (a < 0) {  
  for (i = 0; i < 100; i++) {  
    x[i] = 0  
  }  
}
```

We actually do this in data flow model

```
for (i = 0; i < 100; i++) { // unroll, each element is done separately  
  // one layer (one activation cannot directly influence another activation)  
  y[i] = x[i] + 1000 * a. // (| 1000 * a | >> |x[i]|)  
  z[i] = -x[i] + 1000 * a  
  // second layer ReLu  
  if (y[i] < 0) { y[i] = 0 }  
  if (z[i] < 0) { z[i] = 0 }  
  // third layer  
  xx[i] = y[i] - z[i]  
}
```



# Control Flow Model

We want to do this:

```
if (a < 0) {  
  for (i = 0; i < 100; i++) {  
    x[i] = 0  
  }  
}
```

We actually do this in data flow model

```
for (i = 0; i < 100; i++) { // unroll, each element is done separately  
  // one layer (one activation cannot directly influence another activation)  
  y[i] = x[i] + 1000 * a. // (| 1000 * a | >> |x[i]|)  
  z[i] = -x[i] + 1000 * a  
  // second layer ReLu  
  if (y[i] < 0) { y[i] = 0 }  
  if (z[i] < 0) { z[i] = 0 }  
  // third layer  
  [i] = y[i] - z[i]
```

- Very inefficient (each element is learned independently)
  - Result to large models
- Difficult to understand (decisions are carried out at micro level)
  - Result to interpretation difficulties

# Control Flow Model

We want to do this:

```
if (a < 0) {  
  for (i = 0; i < 100; i++) {  
    x[i] = 0  
  }  
}
```

- Very inefficient (each element is learned independently)
  - Result to large models
- Difficult to understand (decisions are carried out at micro level)
  - Result to interpretation difficulties

We actually do this in data flow model

```
for (i = 0; i < 100; i++) { // unroll, each element is done separately  
  // one layer (one activation cannot directly influence another activation)  
  y[i] = x[i] + 1000 * a. // (| 1000 * a | >> |x[i]|)  
  z[i] = -x[i] + 1000 * a  
  // second layer ReLu  
  if (y[i] < 0) { y[i] = 0 }  
  if (z[i] < 0) { z[i] = 0 }  
  // third layer  
  [i] = y[i] - z[i]
```

How to design a control flow module to

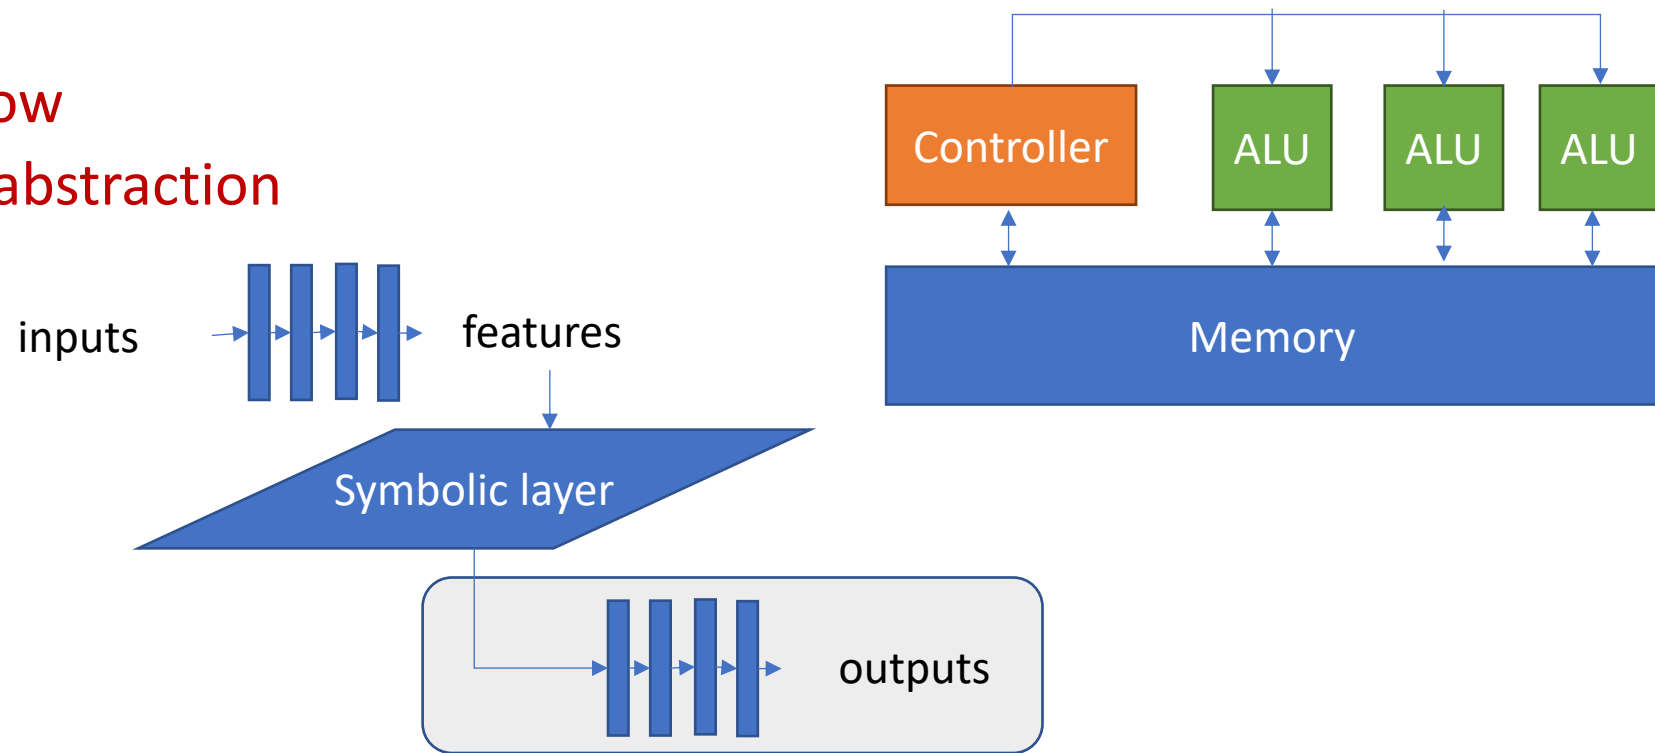
- Reduce the model size
- Better explanation

Challenges

- How does back propagation work?
- What about training stability?

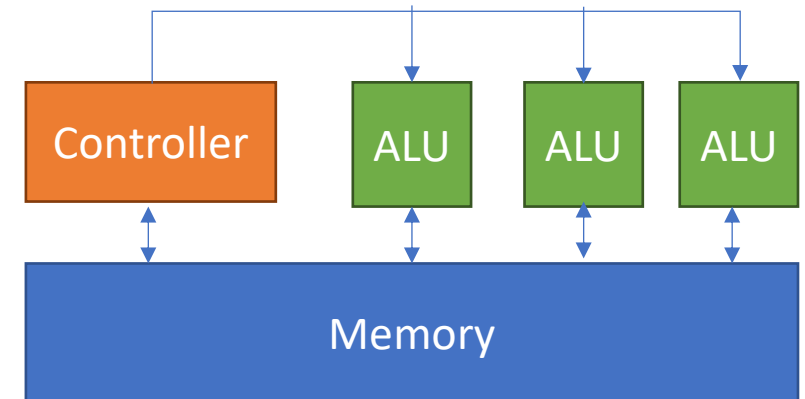
# Trade Time for Space

- May we find a small model, looping inputs many times, and get results of the same quality as the big models?
- Requires
  - Memory
  - Control flow
  - Symbolic abstraction



# Next Generation of Reinforcement Learning

- AI needs to explore via deep reinforcement learning
  - DNN generates heat map
    - Sample actions based on importance
  - RL performs playout
    - Predict the future
- Issues of RL
  - Has fixed actions and fixed states
  - Has a large branching factor
  - Has a local view of the actions
- What's next?
  - Global view, built-in composition
  - DNN generates heat map from control flow



# An analogy of DNN to VLSI

|         | DNN  | VLSI  |
|---------|--|---|
| Stage 1 | A feedforward network with a single layer is sufficient to represent any function  | All combination logic can be represented as two level AND/OR gates  |
| Stage 2 | Deeper networks generalize better. ( <b>what about complexity? Any theory support?</b> )   | Apply multi-level logic optimization to reduce complexity, supported by Boolean algebra   |
| Stage 3 | Softmax (one hot, with prior knowledge) or embedding to coalesce values to finite possibilities. Abstraction? Maybe <b>symbolic layer</b>  | Analog -> digital. Coalesce unlimited analog signals to finite possibilities (0/1). Clean abstraction. Analog: physics, digital: Boolean logic      |
| Stage 4 | Recurrent neural network: <b>Unroll</b> the network in each time step. Perform the <b>same</b> computation in each time step. (Training complexity increases linearly with time steps) | Sequential logic: outputs depend on current and past inputs<br>Finite state machines: backward edge perform <b>different</b> computation each cycle |
| Stage 5 | <b>DNN is purely data flow</b>   | Software directs FSM transition (control flow), producing limitless possibilities   |

# An analogy of DNN to VLSI

|         | DNN   | VLSI  |
|---------|---|---|
| Stage 1 | A feedforward network with a single layer is sufficient to represent any function               | All combination logic can be represented as two level AND/OR gates  |
| Stage 2 | Deeper networks generalize better. ( <b>what about complexity? Any theory support?</b> )        | Apply multi-level logic optimization to reduce complexity, supported by Boolean algebra   |
| Stage 3 | Softmax (one hot, with prior knowledge) or embedding to coalesce values to finite possibilities | Analog -> digital. Coalesce unlimited analog signals to finite possibilities (0/1). Clean abstraction. Analog: physics, digital: Boolean logic      |
| Stage 4 | Control flow to perform the next step. This is linearly separable                               | Sequential logic: outputs depend on current and past inputs<br>Finite state machines: backward edge perform <b>different</b> computation each cycle |
| Stage 5 | <b>DNN is purely data flow</b>  | Software directs FSM transition (control flow), producing limitless possibilities   |

This is where all the magic happens

- We are not “approximating” a function.
- We cannot fully control the outcome.
- This is the most “rewarding” and most “dangerous” step.

# The Fourth AI wave

- Abstract symbolic layer – neurosymbolic computing
- Control flow model
- Revolutionized reinforcement learning
  - Apply DNN based control flow
- Application wise
  - Robotics

A Practical Way to Become  
the Next NVIDIA



# The Unfortunate Truth

- Let's face it. NVIDIA is a giant
- It is unrealistic to hope to build a 10x faster hardware to beat NVIDIA on DNN
- GPU is a good **general** architecture, suitable for workloads that are
  - Data intensive
  - High fine-grained parallelism
  - High throughput
  - Latency insensitive

# A More Practical Way

- Old markets dry up, new markets emerge.
- Let's search a new market
  
- NVIDIA: 2000, 2012
- ARM: 1996

# A More Practical Way

- Find a new market that hit GPU's pain point
  - Control intensive
  - High middle-grained parallelism
  - Latency sensitive
  - High throughput
- Which market?

# The Fourth AI Wave

- Abstract symbolic layer – neurosymbolic computing
  - Requirement of compute on data is not as high
- Control flow model
  - Single batch training – much less fine-grained parallelism
  - Memory – short sequence of random access
  - Control/data interaction – data depends on control, latency sensitive
    - Resource requirement is very high for high latency workload
    - Activation sparsity
- Revolutionized reinforcement learning
  - Enhance RL with DNN based control flow
  - Models will be executed lots of times – compute requirement is high (in a different sense)

GPU is terrible for these workloads

# But wait... There is still Intel/AMD

- CPU is good for workloads that are
  - Control intensive
  - Coarse-grained parallelism
  - Latency sensitive
- But CPU has
  - High context switch cost
  - Too big cores
  - Compute capacity is low
  - Synchronous architecture

# A New Architecture Will Arise

- With new workload, there is a middle ground between CPU and GPU
  - This is the next big opportunity

# A New Architecture Will Arise

- With new workload, there is a middle ground between CPU and GPU
  - This is the next big opportunity
- However, we need to first make sure the market is **real**
- So we need to first work on the algorithm part...

# Research Directions

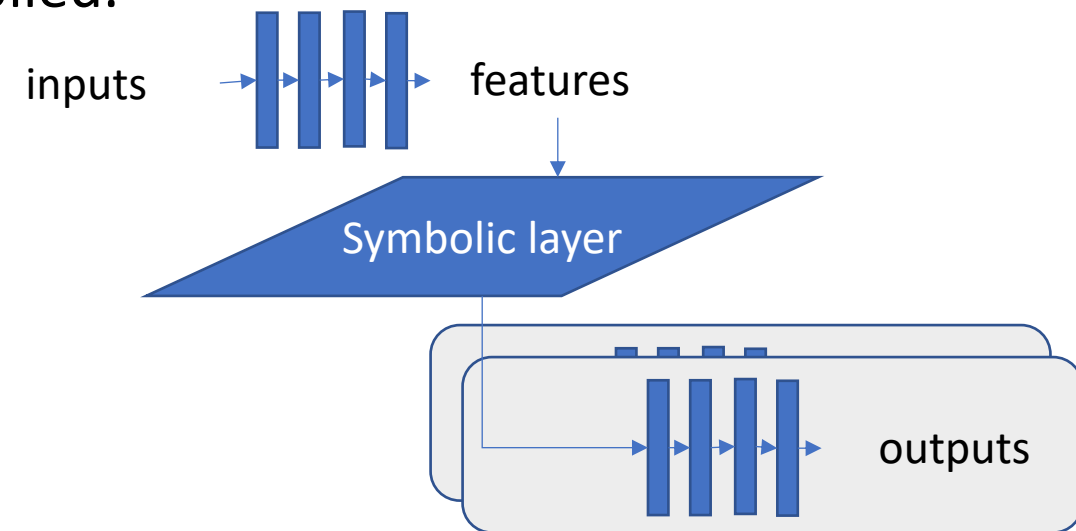


# New Model Architecture Exploration

- DNN: output is inputs' linear combination
- CNN: take advantage of **prior knowledge** and locality
- Self-attention: output is inputs' pair-wise **weighted** linear combination
  
- How to reduce prior knowledge?
  - Research on parameter sharing and sparsity
- Think beyond linear combination?

# Symbolic Layer

- Symbolic layer is for multi-tasking model
  - Symbols are just embeddings for multiple tasks
- Start from connectionist AI and organically grow the symbols
- Use DNN to extract features, abstract the symbolics, and apply DNN on the symbols
  - All done in training, no prior knowledge applied.



# Memory

- Dual purposes:
  - An upgrade of RNN
    - More complex logic between time steps
  - Decouple knowledge and methods
    - Symbolic layer
    - Multi-task model handling
- Approach: start from programs we know about (NTM, DNC, etc.)
  - Impose prior knowledge first and gradually relax the requirement.

# Control Flow

- Activation sparsity
- More complex control logic
  - More than just single layer MLP. Beyond MOE.
- Feedback loop
  - With memory

# Reinforcement Learning

- Familiar with existing reinforcement learning first
  - Identify tough tasks existing solutions fail
- Based on previous researches, apply new findings to reinforcement learning
- E.g. the issue for deep learning + MCTS is that the branching factor is too large.
  - This is a very local view of the actions.
  - Need to increase the steps for the actions, and then decompose the actions to finer grained steps.
    - Do this organically and dynamically without prior knowledge

# But... Where is the Hardware?

- Hardware architecture research follows closely after algorithm breakthroughs
  - Research on algorithm and hardware are pipelined.
  - They require different skills but can work closely together.
- Closely innovate in hardware together with algorithm.
- This gives us 3-4 years lead time for the new hardware to enter market.
  - This is our **best** opportunity to beat NVIDIA/Intel

# Research Approaches

- Option 1: Fast paced iterations on key technologies
  - Pros: technology first, no redundant work, progress fast, small team. May produce dozens of top conference papers.
  - Cons: need extra work to land technology to real applications.
- Options 2: Conquer a few high impact domains with new technology
  - Pros: Better visibility, practical use cases. May produce several nature/science level papers.
  - Cons: Only a small percentage of effort targets new technology. Progress slower. Teams may be bigger. The technology have the risk of being addressed by others.

# Conclusion

- I've talked about my real research objective.
- I've identified the breakthroughs that will lead to the next AI revolution
- I've outlined current research directions that will lead to the breakthroughs



# Purpose of the talk

- I'm looking to **build a small team** with complementing skills and **share the same goal**.
  - DNN model architecture design.
  - Reinforcement learning.
  - ML theory.
  - Computer architecture.
- We dash towards the goal.
- I don't want to cover all aspects, and I cannot.
- But before I find the team, I will try to understand the stake holding directions and consult the rest.

# Purpose of the talk

- If you don't share the same goal and don't agree with my philosophy:
  - I'm looking for collaborators.
- If you feels that some of the proposed directions are of interest to you, let's talk.

# Purpose of the talk

- If you have counter arguments to my story:
- **Let's have a constructive chat.**
  - To help me understand different opinions and improve my thinking.

梦想还是要有的，万一实现了呢？

一群有情有义的人，在一起做一件有价值意义的事

此时此刻，非我莫属

Will you join the venture?