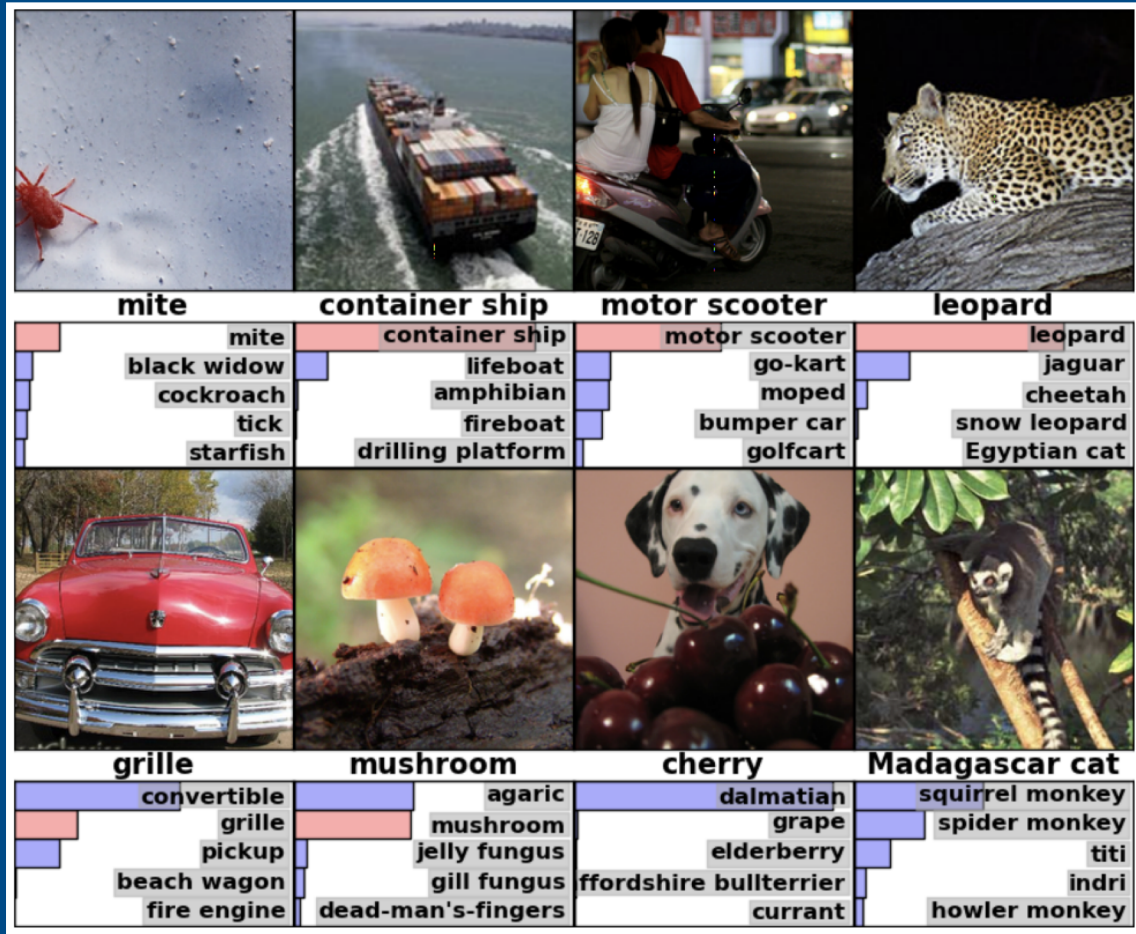


Networking for AI: Challenges and Opportunities

Somdeb Majumdar

Head of AI Lab (US) - Intel Labs



AlexNet. 2012. 60M Parameters.

Title: United Methodists Agree to Historic Split

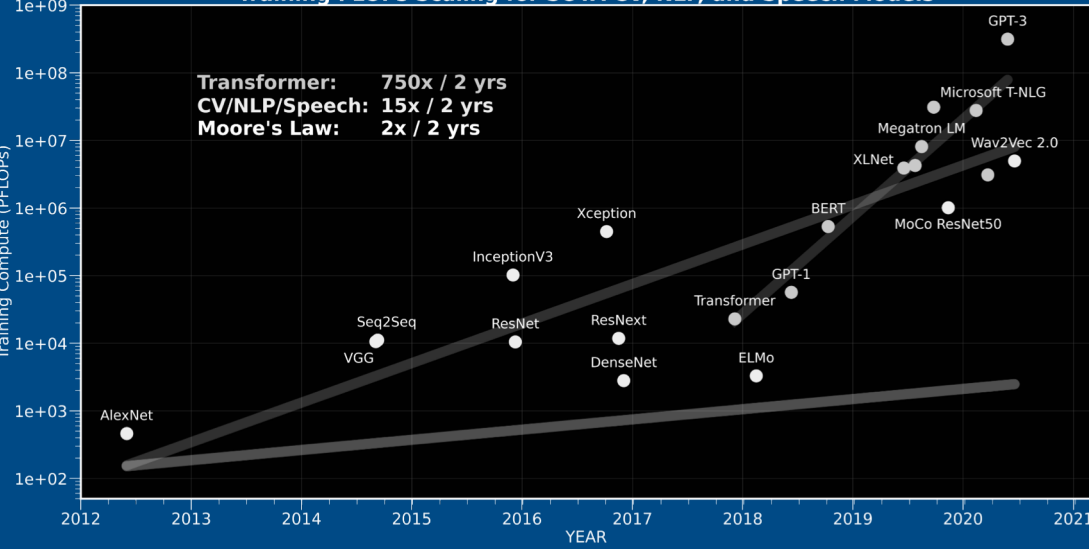
Subtitle: Those who oppose gay marriage will form their own denomination

Article: After two days of intense debate, the United Methodist Church has agreed to a historic split - one that is expected to end in the creation of a new denomination, one that will be "theologically and socially conservative," according to The Washington Post. The majority of delegates attending the church's annual General Conference in May voted to strengthen a ban on the ordination of LGBTQ clergy and to write new rules that will "discipline" clergy who officiate at same-sex weddings. But those who opposed these measures have a new plan: They say they will form a separate denomination by 2020, calling their church the Christian Methodist denomination.

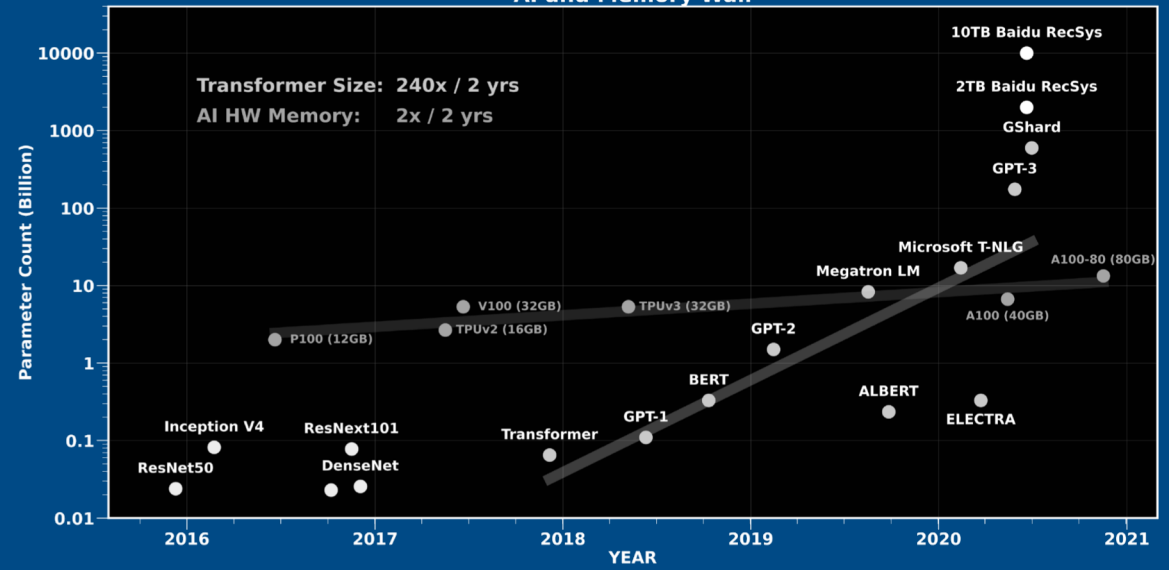
The Post notes that the denomination, which claims 12.5 million members, was in the early 20th century the "largest Protestant denomination in the U.S.," but that it has been shrinking in recent decades. The new split will be the second in the church's history. The first occurred in 1968, when roughly 10 percent of the denomination left to form the Evangelical United Brethren Church. The Post notes that the proposed split "comes at a critical time for the church, which has been losing members for years," which has been "pushed toward the brink of a schism over the role of LGBTQ people in the church." Gay marriage is not the only issue that has divided the church. In 2016, the denomination was split over ordination of transgender clergy, with the North Pacific regional conference voting to ban them from serving as clergy, and the South Pacific regional conference voting to allow them.

GPT3. 2020. 175B parameters.

Training FLOPs Scaling for SOTA CV, NLP, and Speech Models

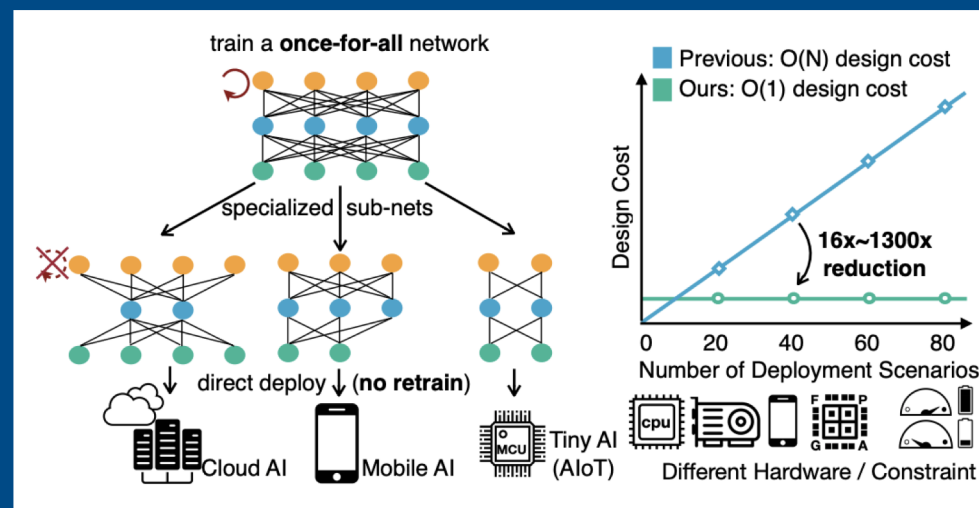


AI and Memory Wall



The Networking Challenges

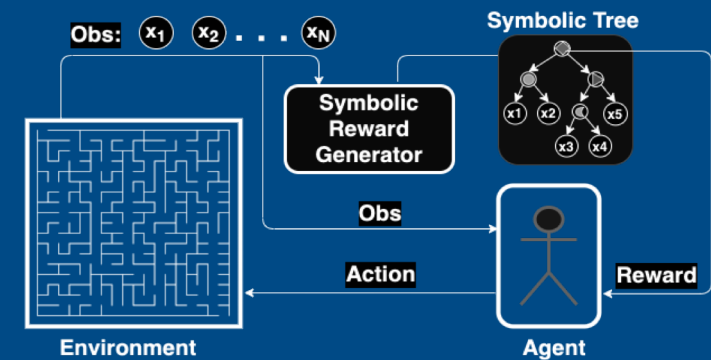
- Distributed training from distributed data
- Memory bottlenecks - especially on edge devices
- Data security
- Can we even train on the edge?
- Hardware-aware models - avoid re-training



[Once-for-All](#), Cai et al., ICLR 2020

The Networking Opportunities

- Multi-agent collaborative systems
- Multi-task learning
- Massive-scale combinatorics
- Homomorphic encryption
- Debuggable AI



```
def get_intrinsic_reward(s_0, s_1, s_2, s_3, s_4, s_5, s_6, s_7):  
    p_1 = tan(cos(s_4)); p_2 = cos(s_3); p_3 = pass_smaller(p_1, p_2)  
    x_1 = multiply(-1, abs(subtract(s_7, p_3)))  
    q_1 = multiply(-1, abs(subtract(1, s_4)))  
    q_2 = max([s_2, 1, s_7, q_1, 0])  
    q_3 = max([q_2, s_7, cos(0), multiply(s_0, s_6), multiply(s_5, subtract(s_6, 1))])  
    y_1 = div_by_10(q_3)  
    y_2 = square(s_7)  
    y_3 = protected_div(1, div_by_100(s_0))  
    x_2 = gate(y_1, y_2, y_3)  
    z = equal_to(x_2, x_1)  
    reward = add(0, pass_smaller(div_by_10(s_7), z))  
    return reward
```

[Learning Intrinsic Symbolic Rewards](#), Sheikh et al, Neurips DRL WS, 2020

Intel AI Lab

- **Who we are: ~50 ML enthusiasts with two core objectives**

- **Understand the fundamental representational properties of modern ML systems**

- **Solve large scale problems using the best of “new” and “old” AI**

- **What do we work on?**

- **All of the challenges and opportunities above - and then some!**

Our Collaborations

