Today's Session Is Being Recorded

This session is being recorded and will be made available for viewing by others by invitation.

If you do not wish for your voice to be recorded, please place your phone on mute for the duration of the session.

This session is non-confidential, so if you wish to ask a question or make a comment, please do not disclose any information that is confidential to you, your employer, or any third party.

The views expressed in this session are those of the contributors and do not necessarily reflect the views of Intel Corporation.

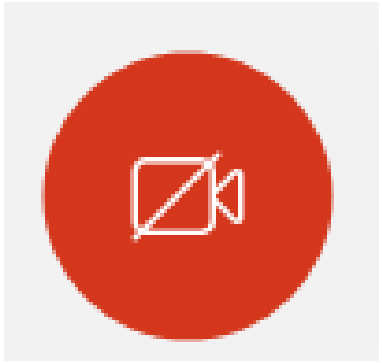

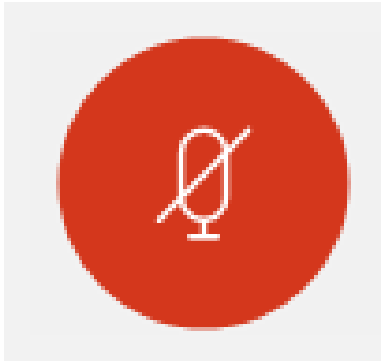

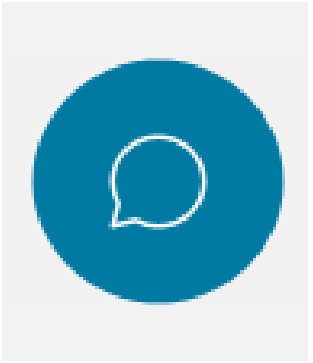

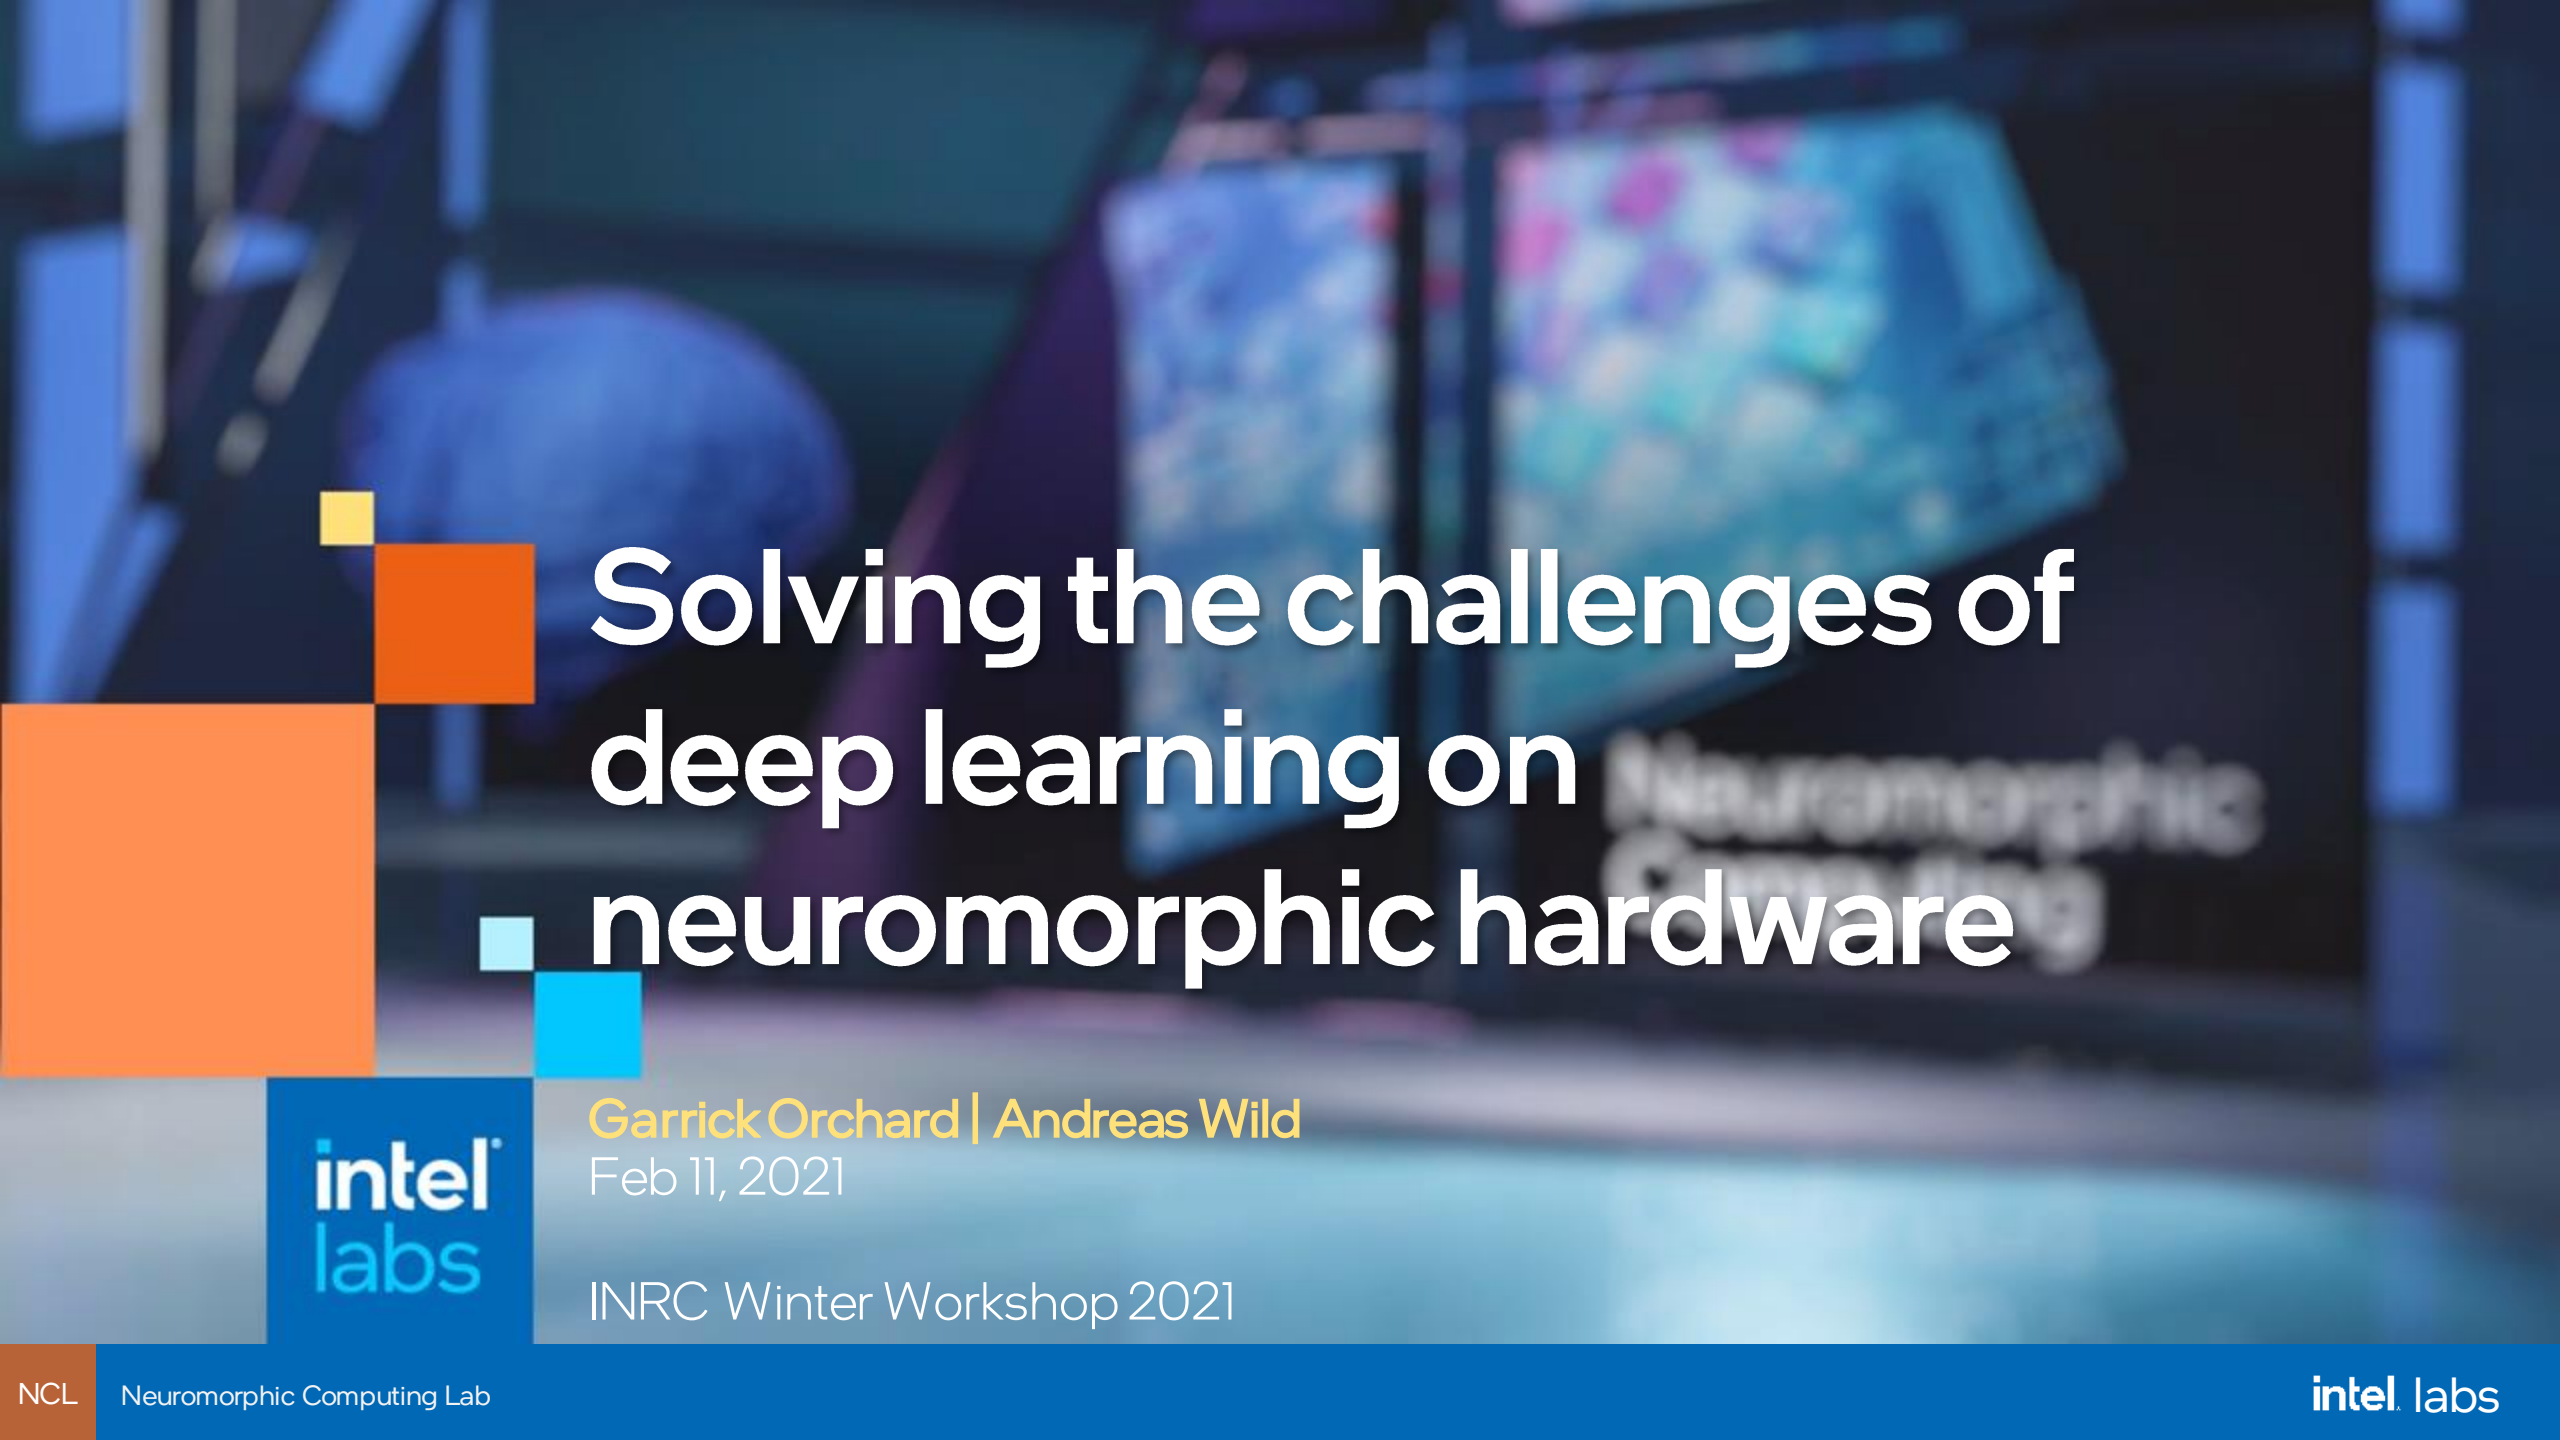

Solving the challenges of deep learning on neuromorphic hardware

Garrick Orchard | Andreas Wild

Feb 11, 2021

INRC Winter Workshop 2021

Legal Information

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Results have been estimated or simulated.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Session goals

- Bring together leaders in the field of neuromorphic DL
- Assess where we are today
- Discuss the major algorithmic challenges that community must address
- Identify key capabilities HW must support

Panelists

- Guillaume Bellec (EPFL)
- Wolfgang Maass (TU Graz)
- Konstantinos Michmizos (Rutgers University)
- Emre Neftci (UC Irvine)
- Priya Panda (Yale)
- Qinru Qiu (Syracuse University)
- Kaushik Roy (Purdue University)
- Walter Senn (University of Bern)
- Friedemann Zenke (Friedrich Miescher Institute Basel)

Format

Learning from this week:

It's easy to run out of time!

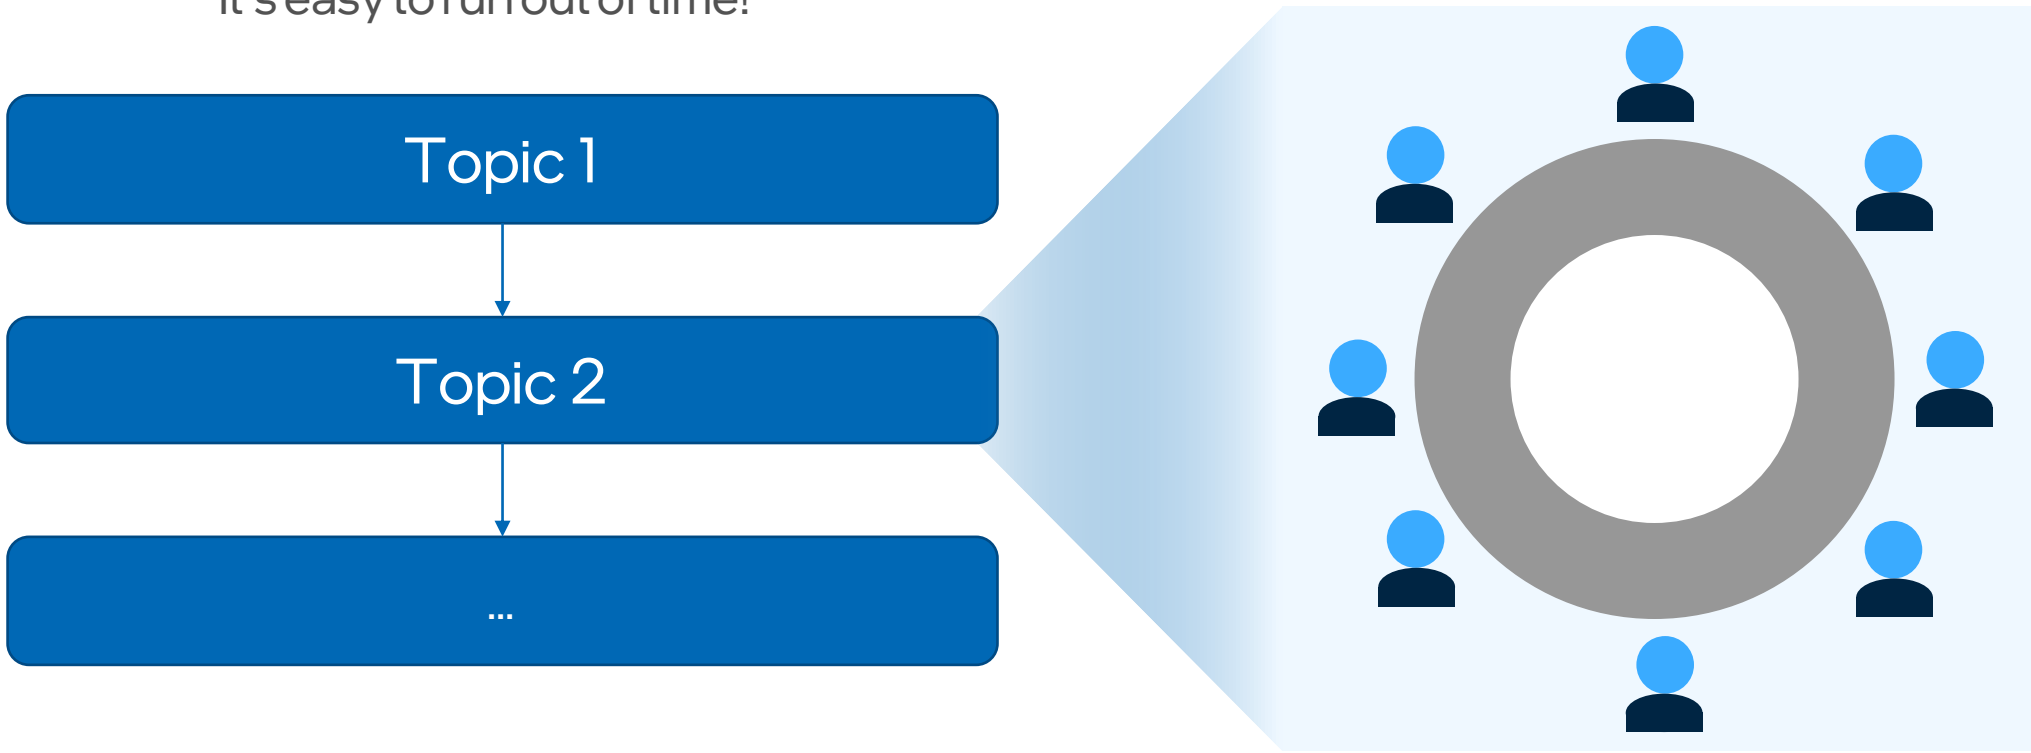

- Not too many topics → broad focus on major areas: offline/online training and inference
- Panelists with an opinion share their perspective one topic at a time and engage in meaningful exchange
- Still keep topics time-bounded to around 10-15min
- Audience may weigh in as well

Topics

- Getting started:

  What's the proper role and meaningful goals for DL in neuromorphic computing?

- Going deeper:
  - Offline training: What are paths towards scalable training of deep SNNs?
  - Online learning: How to achieve the vision of continual online learning from incremental data?
  - Inference: Given deep SNN's strengths and weaknesses, how to best put them to work?

Deep Learning is only part of the SNN algorithm space

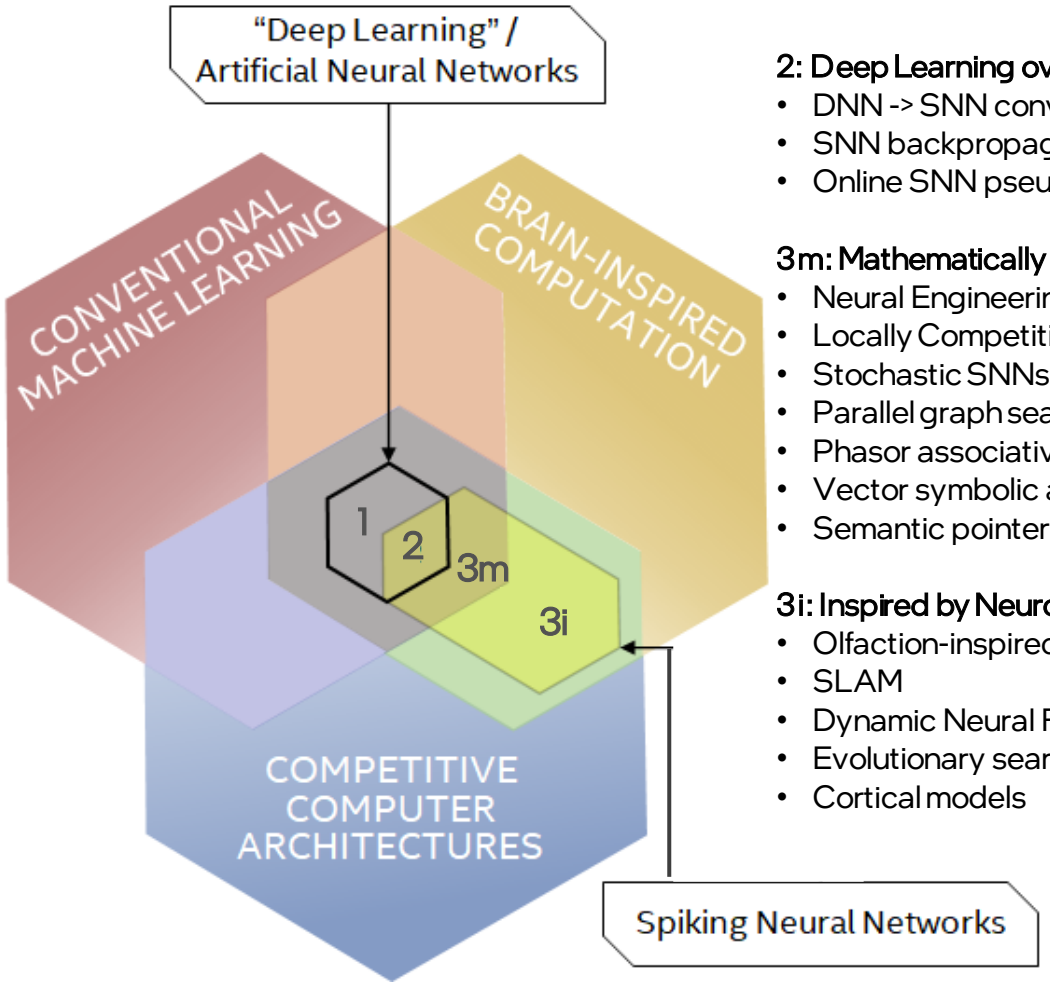

2: Deep Learning overlap with SNNs
- DNN -> SNN conversion
- SNN backpropagation
- Online SNN pseudo-backprop

3m: Mathematically Formulated SNNs
- Neural Engineering Framework (NEF)
- Locally Competitive Algorithm for LASSO
- Stochastic SNNs for solving CSPs
- Parallel graph search
- Phasor associative memories
- Vector symbolic architectures (VSA)
- Semantic pointer architecture (SPA)

3i: Inspired by Neuroscience
- Olfaction-inspired rapid learning
- SLAM
- Dynamic Neural Fields
- Evolutionary search
- Cortical models

What is Deep Learning to SNNs?

- Architectures/Models we should try to implement (DNNs)?
- A tool for configuring SNNs?
- An algorithm we should try implement on-chip/online?

Context

Food for thought to seed & inspire discussion

What's the proper role and meaningful goals for DL in neuromorphic computing?

- Conventional DNNs and many static datasets are not ideal for neuromorphic hardware but are also often the go-to-solution for newcomers into the field.
- DNNs and DL are not the be-all-end-all algorithms, but they can be useful in neuromorphic computing when applied to the right sub-problems.
- Can Loihi provide some new capability to DL?
- What could be the "ImageNet" of neuromorphic computing?

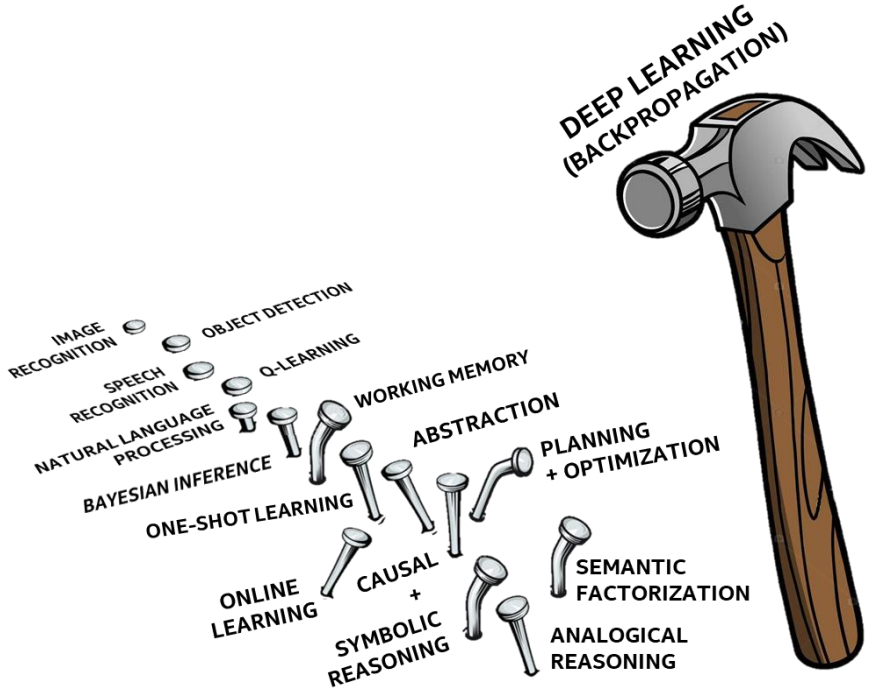

What are paths towards scalable training of deep SNNs?

- Offline SGD training is also useful for deep SNNs even if just for initialization for subsequent online training
- Tools like SLAYER, SpyTorch and other surrogate gradient methods work well to train mostly small deep SNNs so far
- But BPTT is expensive
- How can it be scaled up to large problems?
- What conversion or hybrid training approaches make the most sense?
- Should we pursue the same topologies for SNNs compared to ANNs?
- How do we improve data-efficiency?
- Should there be a stronger focus on ultra-sparse DNNs?

How to achieve the vision of continual online learning from incremental data?

- Online SGD shows little promise on neuromorphic systems:
  - Not a SIMD architecture
  - Limited memory to hold large batches of data for multi-epoch training
- Where do we need online learning versus a sufficiently generalizing offline trained network? When there's not much training data, at the edge, when privacy matters, when models are too large to always retrain from scratch etc....
- Online deep SNN learning only suitable in (real-time) batch=1 regime:
  - Not necessarily a useful approach for dense, statically available data (may as well use GPU)
  - But batch=1 learning is the only choice when data only becomes incrementally available, can't be batched and does not repeat in iid. way
  - New knowledge must be incorporated incrementally and learning must avoid forgetting
- Many 3F learning rules involve continuous post-synaptic state variables. How to ensure weight updates remain sparse to avoid frequent expensive memory updates in neuromorphic HW?
- What are useful tasks to benchmark continual online learning?

Given deep SNN's strengths and weaknesses how to best put them to work?

- What are the best use cases?
- We CAN replicate almost any traditional DNN for inference tasks with SNNs but should we?
- Clearly something in the temporal domain… but let's not just unroll otherwise static data if we don't exploit inhibitory effects to prune activity drastically.
- Do we always have to massively outperform conventional DNNs as part of a larger neuromorphic application where a deep SNN only plays one role within a larger system?

Thank You!