# CPSC 532X: Adaptation & Adaptive Computation

**Who**: Prof. Evan Shelhamer `shelhamer@cs.ubc.ca`
**When**: Tu./Th. 10:00-11:30am 2025W1
**Where**: Orchard 4016

**Note:** This syllabus is still adapting and updating, especially concerning the particular papers and dates for readings, but it is representative of the course content 🚧

## Summary

Machine learning has delivered incredible progress by turning ever more data into models and systems for more and more tasks. However, standard methods learn just once during their training, and then remain the same during their testing or deployment. The course covers how models can keep updating, by adapting what is modeled (adaptation) and how it is executed (adaptive computation), in order to improve accuracy and efficiency. For adaptation we will emphasize test-time adaptation: how to update on new and different data during the deployment of a machine learning system. Such updates may alter the model parameters, inputs, or outputs by further processing. For adaptive computation we will emphasize improvements to computational efficiency and accuracy that alter execution across depth, width, and inputs. For both topics we will focus on improvements that are realized on existing hardware and software, that is, practical methods that improve due to their algorithms and engineering. Join us to update your own knowledge on current topics in AI by discussing papers and completing a course project!

## Schedule

**Office Hours** TODO(shelhamer)

**2025-09-04** Lecture #1 Test-Time Adaptation + Adaptive Computation

**2025-09-09** Lecture #2 Tour of Test-Time Adaptation pt. 2

**2025-09-11** Lecture #3 Tour of Adaptive Computation

**2025-09-16** Live Reading of Tent + Q&A + Lecture #4 1/2

**2025-09-18** Live Reading of TTT + Q&A + Lecture #4 2/2

**2025-09-23** Papers on Test-Time Adaptation (TTA) + Test-Time Training (TTT)

- [Tent] https://arxiv.org/abs/2006.10726 Wang*, Shelhamer* et al. ICLR'21
- [TTT] https://arxiv.org/abs/1909.13231 Sun et al. ICML'20

- bonus: more material and Q&A on adaptation and transfer learning

**2025-09-25** Papers on Episodic and Continual Adaptation

- [MEMO] https://arxiv.org/abs/2110.09506 Zhang et al. NeurIPS'22
- [ETA] https://arxiv.org/abs/2204.02610 Niu*, Wu*, Zhang* et al. ICML'22

**2025-09-30** **no class** UBC is closed for Truth and Reconciliation Day
**2025-10-02** Papers
**2025-10-07** Papers
**2025-10-09** Papers
**2025-10-14** Papers + Project Proposals(!)
**2025-10-16** Papers
**2025-10-21** Papers
**2025-10-23** Papers
**2025-10-28** Papers
**2025-10-30** Papers
**2025-11-04** Papers
**2025-11-06** Papers
**2025-11-11** Papers
**2025-11-13** Papers
**2025-11-18** Papers
**2025-11-20** Papers
**2025-11-25** Papers
**2025-11-27** Papers
**2025-12-02** Projects
**2025-12-04** Projects

## Format

**Lectures** The twin topics of the course, adaptation and adaptive computation, will be introduced by 1+ hour lectures at the beginning of the course. Orientation to additional major topics will be provided as needed in 0.5 hour lectures. 2-3 guest lectures will be given by experts in 0.5 and 1 hour slots.

**Readings, Reactions, and Presentations** We will read 2-3 papers for each class (= 2-6 per week),
discuss in class, and write short reactions for each paper. The reaction is a <256 word summary of your perspective on the paper—for instance what informed, confused, or surprised you—due by email to shelhamer@cs.ubc.ca the night before class (midnight Vancouver time). Each student is responsible for at least one paper presentation: a 30 minute guided discussion of the paper with supporting materials such as slides optional but encouraged.

**Reactions**
what: your perspective on the paper—an optional but suggested template is to include the goal you chose for your reading then what you learned and did not learn to this end
when: due the night before class by midnight Vancouver time
how: Please submit by email to shelhamer@cs.ubc.ca with the subject line "[cs532x] reaction: [PAPER]" where PAPER = the tag given in the schedule like `[Tent]`. Format the reaction as plaintext, markdown, or PDF.

**Presentations**
TODO(shelhamer)

**Course Projects** Every student will complete a technical project on adaptation, adaptive computation, or their intersection. The project may focus on experiment or engineering, but in all cases it must be new and informative, such that the student can identify and explain its contribution(s). For the sake of examples, the implementation and experimentation of a new method qualifies, the reproduction of an existing method with new engineering on a different platform or improved system characteristics qualifies, and the survey of existing methods without original experiments or engineering does not qualify. Course project proposals are due at the midpoint of the class (= Oct. 14) to ensure feasibility and course fit and to provide feedback. Course project poster presentations will be held the last two weeks of class and written project reports are due by the last day of class.

## Grading

20% attendance and participation in class by commenting, questioning, and discussing the material
30% presentations and reactions to the readings
50% course project

## Goals

- **Recognize and understand** key papers on adaptation + adaptive computation and the frontiers of these topics.
- **Read, compare/contrast, and explain** academic papers in machine learning.
- **Present and discuss** technical material on machine learning to inform yourself and others in words, code, visualizations, and so on.
- **Complete a technical project** that is new to you and explain it to others as an exercise in the dissemination of experimental and engineering material.

## Prequisites and Audience

The class material and project expectations require a background in deep learning. Basic knowledge of machine learning and familiarity with deep learning are expected for sufficient understanding of the material and success in the course project. (for example: nonlinearities like the ReLU, convolutional / recurrent / attentional network architectures, optimization by stochastic gradient descent, and competence with a deep learning framework including but not limited to PyTorch or JAX).

This is a graduate-level class. Students pursuing related research in AI are encouraged to join regardless of their department. Note that registered students are expected to complete novel and rigorous course projects that would constitute a strong start to a research publication or open-source engineering project. Undergraduate are welcome with sufficient background and permission of the instructor. Auditors may join provided sufficient capacity.

## References

This is a seminar-style class driven by papers and presentations. The assigned and optional readings are selected from this complete list of publications old and new:

Essential Test-Time Adaptation:

- Updating by Entropy Minimization and the TTA setting: Tent. Wang* and Shelhamer* et al. ICLR’21.
- Updating by Self-Supervision and the TTT setting: TTT by self-supervision. Sun et al. ICML’20.
- Updating Statistics: BN. Schneider et al. NeurIPS’20.
- Updating Outputs: T3A. Iwasawa and Matsuo. NeurIPS’21. LAME. Boudiaf et al. CVPR’22.
- Updating by Diffusion: DiffPure. Nie et al. ICML’22. DDA. Gao* and Zhang* et al. CVPR’23. Diffusion-TTA. Prabhudesai et al. NeurIPS’23.

Essential Adaptive Computation:

- Dynamic neural networks: A survey. Yizeng et al. PAMI'21.
- Zero Time Waste: Recycling Predictions in Early Exit Neural Networks. Wolczyk et al. NeurIPS'21.
- Switch Transformers. Fedus* & Zoph* and Shazeer. JMLR'22.
- From Sparse to Soft Mixtures of Experts. Puigcerver* & Riquelme* et al. ICLR'24.
- Deep Equilibrium Models. Bai et al. NeurIPS'19.
  Neural Deep Equilibrium Solvers. Bai et al. ICLR'22.
  Streaming Multiscale Deep Equilibrium Models. Ertenli et al. ECCV'22.
- Which Tokens to Use? Investigating Token Reduction in Vision Transformers. Haurem et al. ICCVW'23.
- Mixture-of-Depths. Raposo et al. arXiv'24.
- LookWhere? Efficient Visual Recognition by Learning Where to Look and What to See from Self-Supervision. Fuller* & Yassin* et al. arXiv'25.

Multi-Modeling:

- Ensembles. Dietterich. MCS'00.
- Deep Ensembles. Lakshminarayanan et al. NeurIPS'17.
- Model soups. Wortsman et al. ICML'22.
- Seasoning Model Soups. Croce et al. CVPR'23.

Evaluation:

- Evaluating Adaptive Test-Time Defenses. Croce*, Gowal*, Brunner*, Shelhamer* et al. ICML'22.
- Better Practices for Domain Adaptation. Ericsson et al. AutoML'23.

Training and Transfer:

- Fine-tuning and Feature Reuse: DeCAF. Donahue* and Jia* et al. ICML'14.
- How Transferable are Features? Yosinski et al. NeurIPS'14.
- Head2Toe. Evci et al. ICML'22.
- Surgical Fine-Tuning. Lee* and Chen* et al. ICLR'23.
- Parameter Efficient Fine-Tuning: LoRA. Hu* and Shen* et al. ICLR'22. Adapters. Houlsby et al. ICML'19.

More Test-Time Adaptation:

- Survey of TTA. Liang et al. IJCV'25.
- MEMO. Zhang et al. NeurIPS'22.
- EATA. Niu*, Wu*, Zhang* et al. ICML'22.
- CoTTA. Wang et al. CVPR'22.
- NOTE. Gong et al. NeurIPS'22.
- RoTTA. Yuan et al. CVPR'23.
- RDumb. Press et al. NeurIPS'23.
- SAR. Niu*, Wu*, Zhang* et al. ICLR'23.
- PeTTA. Hoang et al. NeurIPS'24.
- FOA. Niu et al. ICML'24.
- GDA. Tsai et al. CVPR'24.

More Test-Time Training:

- TTT+. Liu et al. NeurIPS'21.
- TTT-MAE. Gandelsman* and Sun* et al. NeurIPS'22.
- TTT for RL. Hansen et al. ICLR'21.
- TTT for LLMs. Hardt and Sun. ICLR'24.

More Adaptive Computation:

- BlockDrop. Zuxuan et al. CVPR'18.
- Adaptive computation with elastic input sequence. Fuzhao et al. ICML'23.
- Deep Equilibrium Optical Flow Estimation. Bai et al. CVPR'22.
- Iterative Amortized Inference. Marino et al. ICML'18.
  Resolution adaptive networks for efficient inference. CVPR'20.
- Soft Merging of Experts with Adaptive Routing. Muqeeth et al. TMLR'24.
- Anytime Dense Prediction with Confidence Adaptivity. Liu et al. ICLR'22.

Domain Adaptation and Domain Generalization

- Adapting visual category models to new domains. Saenko et al. ECCV'10.
- Online domain adaptation of a pre-trained cascade of classifiers. Jain & Learned-Miller. CVPR'11.
- Survey of Domain Generalization. Zhou et al. PAMI'23.
- Residual Adapters. Rebuffi et al. NeurIPS'17.
- DSI. Yu et al. CVPR'23.