Welcome to Hardware for Machine Learning

Mark Vesterbacka

2020-01-27

Program

1.	Introduction to AI/ML *Fredrik Heintz, IDA*	27/1 *10:15*
2.	Hardware Accelerators and Programming for ML *Christoph Kessler, IDA*	3/2 *10:15*
3.	GPU Architecture and Architecture Dependent Algorithms *Ingemar Ragnemalm, ISY*	10/2 *10:15*
4.	FPGA for ML *Fredrik Medley, Veonner*	17/2 *10:15*
5.	ASIC for ML *Anders Lloyd, Axis Communications*	24/2 *13:15*
6.	Near-Sensor Image Processing *Jörgen Ahlberg, ISY*	2/3 *TBD*
7.	Spiking Networks *Robert Forchheimer, ISY*	9/3 *10:15*
8.	Emerging technologies *Mark Vesterbacka, ISY*	16/3 *10:15*

Assignments

- Assignments are distributed and collected via e-mails with *Mark Vesterbacka*
- Hand-ins should be in *pdf* format
- Assignment schedule:

	Lecture	*Hand-out*	*Hand-in*
1.	Introduction to AI/ML	29/1	7/2
2.	Hardware Accelerators and Programming for ML	5/2	14/2
3.	GPU Architecture and Architecture Dependent Alg.	12/2	21/2
4.	FPGA for ML	19/2	28/2
5.	ASIC for ML	26/2	6/3
6.	Near-Sensor Image Processing	4/3	13/2
7.	Spiking Networks	11/3	20/3
8.	Emerging technologies	18/3	26/3

Grading

- Full Course consists of eight lectures and eight assignments
  - Reward: 6 credits
- Lite Course consists of eight lectures only
  - Reward: 3 credits
- A Missed Lecture is remedied by reading a corresponding journal article
  - Find an article with similar content, read it, and write a one page report on it
- An Unsatisfactory Solved assignment gets one extra chance
  - After this you will be downgraded to the Lite Course

Bibliography

- Great introduction
  - V. Sze, Y. H. Chen, T. J. Yang, J. S. Emer, ”Efficient processing of deep neural networks: A tutorial and survey”, *Proc. IEEE*, vol. 105, no. 12, pp. 2295-2329, Dec. 2017