1. **Course Number and Course Title:**
   COE59406 Hardware Architectures for Machine Learning

2. **Credits Hours:**
   3 – 0 – 3

3. **Prerequisites and/or Co-Requisites:**
   Approval of Program Director and completion of an undergraduate course in Computer Architecture

4. **Name and Contact Information of Instructor:**
   Dr. Assim Sagahyroon

5. **Course Description (Catalog Description):**

   Covers basic machine learning (ML) algorithms; introduces the design and implementation of hardware architectures for efficient processing of ML algorithms, topics include ML on programmable platforms, deep learning accelerators, design alternatives and optimization for ML algorithms, training for deep learning; discusses the characteristics, advantages and disadvantages of popular hardware platforms including CPU, GPU , FPGA in supporting the design of efficient learning networks.

6. **Textbook and other Supplemental Material:**
   Textbook:
   The instructor will provide most of the background material for this course in addition to the use of references listed below.

   References:
   - P. Whatmough, G.Y Wei, D. Brooks, *Deep Learning for Computer Architects*, Morgan and Claypool, 2017
   - M. Vladutiu, Computer Arithmetic: Algorithms and Hardware Implementations, Springer, 2012
   - N. Hemsoth and T. Prickett Morgan, *FPGA Frontiers: New Applications in Reconfigurable Computing,* NextPlatform Press, 2017.
   - C. Kachris, B. Falsafi, D. Soudris, *Hardware Accelerators in Datacenters, Springer*, 2019

7. Learning Outcomes:

Upon completion of the course, students will be able to:

1. Explain the basics of machine learning and its applications
2. Analyze tradeoffs between various hardware architectures and platforms including CPUs, GPUs and FPGAs in running machine learning (ML) applications
3. Identify the key design considerations for efficient ML processing
4. Assess the utility of various design techniques for efficient processing of machine learning.
5. Evaluate future implementation trends and opportunities in deep learning systems.

8. Teaching and Learning Methodologies:

The course will involve a mix of lectures interspersed with paper reading and discussions. A term paper and a semester long project (broken into mini projects) will focus on developing a hardware accelerator for ML, and prototyping it on a FPGA.

9. Course Topics and Schedule:

Topic	Weeks
Introduction to machine learning –Part I	1
Introduction to machine learning –Part II	1
Application examples of machine learning	1
Arithmetic circuits for machine learning processing	1
Available hardware platforms and their characteristics: GPUs	1
Available hardware platforms and their characteristics: FPGAs and ASICs	1
Training versus inference : differences	1
Training versus inference : computational needs	1
Hardware accelerators and accelerated systems architecture	1
Energy efficient acceleration	1
Machine learning on Specialized Hardware	1
Metrics for evaluating machine Learning Systems	1
Use in datacenters	1
Discussion on emerging trends in Machine Learning hardware architectures	1
Presentations	1
Review and Examinations	1
Total:	16