

SM3750: MACHINE LEARNING FOR ARTISTS

Effective Term

Semester A 2022/23

Part I Course Overview

Course Title

Machine Learning for Artists

Subject Code

SM - School of Creative Media

Course Number

3750

Academic Unit

School of Creative Media (SM)

College/School

School of Creative Media (SM)

Course Duration

One Semester

Credit Units

3

Level

B1, B2, B3, B4 - Bachelor's Degree

Medium of Instruction

English

Medium of Assessment

English

Prerequisites

SM2715 Creative Coding

Precursors

Nil

Equivalent Courses

Nil

Exclusive Courses

Nil

Part II Course Details

Abstract

Machine learning pervades many aspects of contemporary life. In response to this situation, media artists have also started to apply and reflect on machine learning algorithms in their own work. This course introduces basic concepts of machine

learning for artists in a hands-on practical way. The focus is not on a rigorous presentation of technical material but on the use of techniques for creative purposes.

The course will have two parts. First, we will introduce the fundamental concept of machine learning and other related ideas (supervised vs. unsupervised learning, regression vs. classification, etc.) and explore classical algorithms in such areas as clustering, classification, dimensionality reduction, and manifold learning. Instead of jumping directly to advanced topics like deep neural networks, we therefore begin with classical algorithms and fundamental notions to build a strong foundation. The second part will then move on to neural networks and deep learning. Students will learn how to use pretrained models and to design simple networks to perform such tasks as image classification, object recognition, semantic segmentation, depth estimation, etc.

The course will mainly concentrate on practical techniques that artists can use. The focus of the course will be on image processing rather than natural language or sound, but students can develop projects in those areas on the basis of the concepts learnt in class. Students are expected to write their own code in Python and to reflect on the techniques that they use from technical, aesthetic, cultural and social standpoints. These aspects can include, for instance: a description of the history of neural networks and machine learning in relation to eugenics, cybernetics, or warfare; a discussion of the social impact of machine learning on gender, work, poverty, or race; a reflection on the philosophical aspects of machine learning, such as the nature of induction; a discussion of political and social aspects of ImageNet and other popular datasets; an awareness of the tendency of technology to become a black box and the problem of interpretable or explainable AI; questions of resource-use (for instance energy consumption, carbon emission, or impact on climate); data collection, digital labour, and surveillance capitalism; etc.

Assessment will be studio-based. Workshops will be conducted using contemporary languages and frameworks, such as for instance: Python, and the main libraries are scikit-learn, scikit-image, Pytorch, or TensorFlow/Keras. The specific languages to be used will depend on the instructor.

Course Intended Learning Outcomes (CILOs)

CILOs	Weighting (if app.)	DEC-A1	DEC-A2	DEC-A3
1 Describe the main machine learning concepts, including regression, classification, clustering, dimensionality reduction, manifold learning, neural networks, and deep learning		x	x	
2 Apply the concepts and techniques in the creation of artistic projects using Python		x	x	x
3 Reflect on the social, cultural, and artistic role or impact of machine learning technologies		x		
4 Articulate one's individual artistic vision regarding the use of machine learning method		x	x	

A1: Attitude

Develop an attitude of discovery/innovation/creativity, as demonstrated by students possessing a strong sense of curiosity, asking questions actively, challenging assumptions or engaging in inquiry together with teachers.

A2: Ability

Develop the ability/skill needed to discover/innovate/create, as demonstrated by students possessing critical thinking skills to assess ideas, acquiring research skills, synthesizing knowledge across disciplines or applying academic knowledge to real-life problems.

A3: Accomplishments

Demonstrate accomplishment of discovery/innovation/creativity through producing /constructing creative works/new artefacts, effective solutions to real-life problems or new processes.

Teaching and Learning Activities (TLAs)

TLAs	Brief Description	CILO No.	Hours/week (if applicable)
1	Lecture	Lectures introducing basic concepts in machine learning, their application in art, and the social or cultural role of machine learning technologies.	1, 2, 3
2	Workshop	Hands-on programming exercises using Python to build familiarity with machine learning techniques.	1, 2
3	Presentation	Description and justification of one's own artistic projects, with reference to aspects of the historical, social or philosophical aspects of machine learning.	1, 3, 4
4	Critique	Participation in critique sessions discussing one another's artistic projects and ideas.	1, 3, 4

Assessment Tasks / Activities (ATs)

ATs	CILO No.	Weighting (%)	Remarks (e.g. Parameter for GenAI use)
1	Presentations + critique	1, 2, 3, 4	40
2	Artistic projects	1, 2, 3, 4	60

Continuous Assessment (%)

100

Examination (%)

0

Assessment Rubrics (AR)**Assessment Task**

1. Artistic projects

Criterion

Ability to apply machine learning concepts and techniques to the production of artworks that show a personal reflection on the cultural aspects of technology.

Excellent (A+, A, A-)

- Strong understanding of the technology being used through the ability to design innovative techniques and wrote one's own code.
- Highly original artistic vision and professional artistic craft.
- Innovative reflection on the cultural and social aspects of technology supported by solid and thorough research.

Good (B+, B, B-)

- Good understanding of machine learning technologies and concepts.
- Expression of a personal artistic vision that is somewhat original and is professionally executed.
- Reflection on social and cultural aspects of machine learning without substantial originality but backed by solid and thorough research.

Fair (C+, C, C-)

- Basic understanding of machine learning technologies using algorithms created by other people without substantial modification.
- Ability to complete artistic projects that lack originality and demonstrate poor artistic craft.
- Basic understanding of the social or cultural impact of technology without any in-depth reflection and limited research.

Marginal (D)

- Sufficient ability to apply the technology in artistic projection to pass the course, often using only code written by other people.
- Lack of artistic originality and deep reflection on social and cultural aspects.
- Below professional artistic standards.

Failure (F)

Inability to understand the technology at any level and apply it to the production of artistic projects.

Assessment Task

2. Presentation and critique

Criterion

Expression of the artistic motivation and structure of the project and its relation to social or philosophical aspects of machine learning.

Excellent (A+, A, A-)

- Clear and accurate description of the technology and the justification, principles, and historical background that underpin the student's art projects.
- Presentations and feedback to other students show an excellent understanding of the technology, an innovative use of the technology, and an innovative artistic project grounded on the theory and/or history of art, informed by original insights into the social or philosophical aspects of machine learning.

Good (B+, B, B-)

- Good description of the technology and the artistic principles and history that underpin the work, with some innovative aspects.
- Clear description of the work in terms of both its technological value and its artistic value, without any substantial technical or artistic innovation, and without innovative insights into the social or philosophical aspects of machine learning.

Fair (C+, C, C-)

- Basic description of the technology but with occasional mistakes or with a superficial understanding of technical issues.
- Reasonable application of this technology to artistic production, without an original artistic vision and with a superficial discussion of art history.

Marginal (D)

- Some understanding of the technology, either superficial or with mistakes, for instance using existing libraries without any innovation.
- Unoriginal artistic project, demonstrating a lack of in-depth reflection on the aesthetic, cultural, or social aspects of the work.

Failure (F)

Inability to describe the technology or the artistic vision or its aesthetic or cultural history.

Additional Information for AR

All A+/A/A- grade assignment should comply with the highest performance of Discovery-oriented learning.

Part III Other Information**Keyword Syllabus**

The following outline is only indicative. Different faculty members who teach this course will highlight different aspects of machine learning.

- a. Introduction to machine learning: fundamental concepts.
- b. Classical machine learning methods using scikit-learn: clustering, linear and non-linear regression, classification, decision trees, dimensionality reduction (PCA, matrix factorization, dictionary learning), manifold learning, etc.
- c. Basic overview of computer vision using scikit-image.
- d. Introduction to neural networks and their history.
- e. Deep learning methods: use of pretrained models, transfer learning, convolutional networks, creating one's own model, encoder-decoder and autoencoder models, working with datasets, self-supervised learning, LSTM networks, transformers.
- f. Cultural, historical, social aspects of machine learning.

These various aspects will not be covered sequentially. For instance, the last point (cultural, historical and social aspects) will be discussed throughout the course.

Reading List**Compulsory Readings**

	Title
1	Scikit-learn user guide and tutorials https://scikit-learn.org/stable/
2	Scikit-image documentation https://scikit-image.org/
3	PyTorch documentation and tutorials https://pytorch.org/tutorials/
4	Machine Learning for Artists https://ml4a.net/
5	Keras guides and examples https://keras.io/guides/
6	AI Art Gallery http://www.aiartonline.com/
7	Algorithmic Justice League https://www.ajl.org/about

Additional Readings

	Title
1	Ethem Alpaydın. Introduction to Machine Learning. MIT Press, 2014.
2	John D. Kelleher. Deep Learning. MIT Press, 2019.
3	Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong. Mathematics for Machine Learning (Cambridge, 2020).
4	Aurélien Géron, Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly Media, 2017.
5	Ian Goodfellow, Yoshua Bengio and Aaron Courville. Deep Learning. MIT Press 2016. https://www.deeplearningbook.org/
6	Hastie, T., R. Tibshirani, and J. H. Friedman. The Elements of Statistical Learning: Data Mining, Inference and Prediction. New York: Springer, 2001
7	Gavin Hackeling. Mastering Machine Learning With Scikit-Learn. Birmingham, UK: Packt, 2014.

8	Gareth James, Daniela Witten, Trevor Hastie and Robert Tibshirani. Introduction to Statistical Learning. Springer, 2013. Corrected 8th printing, 2017.
9	Mitchell, Tom. Machine Learning. New York, NY: McGraw-Hill, 1997.
10	Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press 2012.
11	Francois Chollet. Deep Learning with Python. New York: Manning, 2022.
12	Denis Rothman. Hands-On Explainable AI (XAI) with Python#: Interpret, Visualize, Explain, and Integrate Reliable AI for Fair, Secure, and Trustworthy AI Apps. Birmingham, UK: Packt Publishing, 2020.
13	Shai Shalev-Shwartz and Shai Ben-David. Understanding Machine Learning: From Theory to Algorithms (Cambridge, 2014)
14	Eli Stevens, Luca Antiga, and Thomas Viehman. Deep Learning with PyTorch. New York: Manning, 2020.
15	Ian Pointer. Programming PyTorch for Deep Learning : Creating and Deploying Deep Learning Applications. O’ Reilly, 2019.
16	Chris Mattman. Machine Learning with TensorFlow. New York: Manning, 2021.
17	Natural Language Toolkit documentation https://www.nltk.org/
18	Librosa documentation https://librosa.org/doc/latest/index.html
19	Hugging Face resources and online course about Transformers https://github.com/huggingface
20	Yoav Goldberg. Neural Network Methods for Natural Language Processing. San Rafael, California: Morgan and Claypool, 2017.
21	Hobson Lane, Cole Howard, Hannes Hapke. Natural Language Processing in Action: Understanding, Analyzing, and Generating Text with Python. London: Manning, 2019.
22	Uday Kamath, John Liu, and James Whitaker. Deep Learning for NLP and Speech Recognition. New York: Springer, 2019.
23	Joanna Zylinska. AI Art: Machine Visions and Warped Dreams. Open Humanities Press, 2020. https://openhumanitiespress.org/books/titles/ai-art/
24	Sofian Audry. Art in the Age of Machine Learning. MIT Press, 2021.
25	Ploin, A., Eynon, R., Hjorth I. and Osborne, M. AAI and the Arts: How Machine Learning is Changing Artistic Work. Report from the Creative Algorithmic Intelligence Research Project. Oxford: Oxford Internet Institute, 2022.
26	Cathy O’ Neill. Weapons of Math Destruction. Penguin, 2016.
27	Art and Machine Learning https://medium.com/artists-and-machine-intelligence
28	Schwartzman, Madeline. See Yourself Sensing – Redefining Human Perception. Black Dog Publishing 2011.
29	Massimo Airoidi. Machine Habitus: Toward a Society of Algorithms. London: Polity, 2022.
30	Adam Rutherford. Control. London: W&N, 2022.
31	Adrian Mackenzie. Machine Learners. MIT Press, 2017.
32	Penousal Machado, Juan Romero, and Gary Greenfield, eds. Artificial Intelligence and the Arts. New York: Springer, 2021.
33	Marcus Du Sautoy. The Creativity Code – Art and Innovation in the Age of AI. Harvard University Press, 2019.
34	Jonathan Roberge and Michael Castelle, eds. The Cultural Life of Machine Learning: An Incursion into Critical AI Studies. New York: Springer, 2021.
35	Stegka Hristova, Jennifer Daryl Slack and Soonkwan Hong, eds. Algorithmic Culture: How Big Data and Artificial Intelligence are Transforming Everyday Life. London: Lexington Books, 2020.
36	Soon Yau Cheong. Hands-On Image Generation with TensorFlow. Birmingham, Packt, 2020.
37	Ben Vickers and K. Allado-McDowell. Atlas of Anomalous AI. Ignota Books, 2021.
38	Dreyfus, Hubert. What Computers Can’ t Do. New York: MIT Press, 1972.
39	David Belinksi. The Advent of the Algorithm. Harcourt Books, 200.

40	Roberto Cordeschi. <i>The Discovery of the Artificial: Behavior, Mind and Machines Before and Behind Cybernetics</i> . Springer, 2002.
41	Pedro Domingos. <i>The Master Algorithm</i> . New York: Basic Books, 2015.
42	Jean-Pierre Dupuy. <i>The Mechanization of the Mind: On the Origins of Cognitive Science</i> . Princeton, 2001.
43	Emily Denton, Alex Hanna, and Razvan Amironesei. <i>On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet</i> . <i>Big Data and Society</i> . July 2021.
44	Susan Leavy. <i>Gender bias in artificial intelligence: The need for diversity and gender theory in machine learning</i> . In: <i>Proceedings—international conference on software engineering</i> . IEEE Computer Society, 2018, pp. 14–16
45	Clementine Collett, Gina Neff, and Livia Gouvea Gomes. <i>The Effects of AI on the Working Lives of Women</i> . Report of the Inter American Development Bank Published in 2022 by the United Nations Educational, Scientific and Cultural Organization (UNESCO).
46	Wendy Hui Kyung Chun. <i>Programmed Visions: Software and Memory</i> . MIT Press, 2013.
47	Paul Thagard. “Philosophy and Machine Learning.” <i>Canadian Journal of Philosophy</i> , vol. 20, no. 2, 1990, pp. 261–76.
48	Frank Pasquale. <i>The Black Box Society: The Secret Algorithms That Control Money and Information</i> . Harvard, 2016.
49	Will Douglas Heaven. <i>Predictive policing algorithms are racist. They need to be dismantled</i> . <i>MIT Technology Review</i> . July 17, 2020.
50	Shoshana Zuboff. <i>The Age of Surveillance Capitalism</i> . New York: Public Affairs, 2019.