HCI 560X - Learn to Speak AI

Spring 2021
All sections taught online.
Iowa State University
Ames, Iowa 50011

Instructor: Vladimir Sukhoy

Office:	Howe Hall, Room 2624E
Phone:	630-445-1332 (email preferred)
Email:	sukhoy@iastate.edu
Web Page:	`https://www.ece.iastate.edu/~sukhoy/`
Streaming Lectures:	4:10 – 5:25 p.m. on Tuesdays and Thursdays (optional attendance)
Office Hours:	5:30 – 6:30 p.m. after each streaming lecture, or by appointment

Delivery: online and via recorded videos of lectures.

Course Description: This class covers the main concepts from the design and analysis of algorithms in Artificial Intelligence. We will quickly recap the necessary mathematical knowledge and use it to learn how several popular learning algorithms work. These examples will help illustrate technical issues that specialists in these fields have to address. You will work in small groups to compare the performance of two or more standard learning algorithms using a data set of your choice and describe the results using the material covered in the lectures.

At the end of this class, you will be able to understand technical articles that focus on evaluating standard learning algorithms. You will also have the skills to run standard algorithms on data sets and communicate the results of your experiments to technical audiences that include engineers and computer scientists. Finally, you will acquire an understanding of problems that learning algorithms were designed to solve and will also be able to identify tasks that may not be suitable for these methods.

This course is designed to teach you to be a well-informed user of existing artificial intelligence algorithms. It will not focus on the development of new learning algorithms. There will be no in-depth analysis of convergence, mathematical optimization, choices of programming languages and platforms, etc. The theoretical background will be restricted to the minimum set of concepts that are necessary to understand the problems solved by learning algorithms and the issues that arise due to the need to compute their solutions. The course does require basic computer literacy, i.e., the ability to find information on the Internet, download and install new software, and the ability to work with spreadsheets, e.g., Microsoft Excel.

The course can satisfy the HCI Implementation requirement or serve as an elective that does not. There are assignment options you will have in the course where you can choose between problems designed for programmers or non-programmers. If you want this course to satisfy the program's Implementation requirement, then you must get at least 50% of the credit assigned to all programming problems in all homeworks (i.e., solve at least one half of all problems that have the corresponding label) and also earn a B or better letter grade for the course. Note that this course will not teach you how to program; you must either already have those skills or gain them on your own (e.g., through LinkedIn Learning or elsewhere).

Topics to be Covered: Discrete and continuous probability distributions. Markov models. n-gram models. Underfitting and overfitting. The curse of dimensionality. The computational complexity of algorithms. Supervised learning. Unsupervised learning. Reinforcement learning. Reproducibility of learning algorithms.

Classifiers: k-nearest-neighbors, artificial neural networks, support vector machines.
Regression models: polynomial curve fitting, k-nearest-neighbors regression.
Clustering algorithms: hierarchical clustering, k-means, and Gaussian mixtures (expectation minimization).
Dimensionality reduction: principal component analysis, feature hashing.
Artificial neural networks: the back-propagation algorithm, deep learning.
Evaluation techniques: precision and recall, receiver operating characteristic (ROC) curves, cross-validation.

Readings: The textbook for this class is "Introduction to Machine Learning" by Ethem Alpaydin, 4th edition (3rd and 2nd editions are also OK). In addition to the textbook, the lectures will be based on other sources, most of which can be downloaded from the Internet. All readings outside the required textbooks will be made available on the Canvas page for this course.

Organization and Format*: The course will be organized as a seminar. The expectation is that the students read all material assigned to each lecture before the meeting so that they are ready to discuss it during the meeting. Also, there will be four homework assignments and a course project.

* The instructor can modify any aspect of this class in any way and for whatever reason or for no reason.

Prerequisites: Graduate and advanced undergraduate students can take this graduate class. It is recommended that the students have been exposed to at least 2 of the following fields: linear algebra, statistics, computational perception, human computer interaction, computer vision, signal processing, software engineering, streaming audio or video, and web development.

Programming skills are useful but **not essential**. Data analysis skills with spreadsheet software (e.g., Microsoft Excel) are recommended, but not required because the necessary material will be introduced during the lectures. A subscription to Microsoft Office, which includes Microsoft Excel, is available to all Iowa State students free of charge. Basic computer literacy, i.e., the ability to install and run new software on your machine, download and upload files, etc. is an **essential prerequisite** for this course. Another crucial requirement is your interest in learning algorithms and the desire to understand how they work.

You **will need** programming skills to solve the problems that contribute to satisfying the HCI Implementation requirement.

Required Technology:

1. A reliable high-speed Internet connection (for online students) or regular and dependable access to Iowa State University computer system (for on-campus students).
2. For online students only: access to a computer with a microphone and audio capability.
3. All students: reliable and regular access to a machine that can run Parallels Desktop, VirtualBox, VMWare Workstation, VMWare Workstation Player, VMWare Fusion, or another modern x86-64 virtualization platform that supports a 64-bit guest system distributed in the standard OVF format. The virtual machine used in class requires about 1 gigabyte of RAM and 20 gigabytes disk space. These requirements are only for the virtual machine itself. Your system will need additional RAM and disk space for to run everything else. **Most modern desktops and laptops satisfy this requirement.**

   This technology is required for running the implementations of machine learning algorithms were packaged for this course in a Linux virtual machine. This approach enables exposing students with basic computer literacy to learning algorithms without requiring an IT or computer programming background to set up their implementations. The VM will be able to process only relatively small data sets, i.e., no more than several tens of millions of numerical features across all data instances. Its coverage will be restricted to data sets that have a fixed number of numerical attributes. Text, images, and other types of non-numerical data will need to be converted into numerical form to make the data suitable for the course-provided software. Detailed instructions for using the VM will be provided during the first week of the course.

Students who intend to use Apple silicon Macs for this course: **contact the instructor**.

Free Expression: Iowa State University supports and upholds the First Amendment protection of freedom of speech and the principle of academic freedom in order to foster a learning environment where open inquiry and the vigorous debate of a diversity of ideas are encouraged. Students will not be penalized for the content or viewpoints of their speech as long as student expression in a class context is germane to the subject matter of the class and conveyed in an appropriate manner.

Academic Dishonesty: cheating, plagiarism, and other academic misconducts will not be tolerated and will behandled according to the ISU's academic dishonesty procedures, which are posted here: `http://catalog.iastate.edu/academic_conduct/#academicdishonestytext`

Accommodations: Iowa State University complies with the Americans with Disabilities Act and Sect 504 of the Rehabilitation Act. If you have a disability and anticipate needing accommodations in this course, please contact your instructor to set up a meeting within the first two weeks of the semester or as soon as you become aware of your need. Before meeting with your instructor, you will need to obtain a SAAR form with recommendations for accommodations from the Student Accessibility Services (`http://new.dso.iastate.edu/dr/student`) located in Room 1076 on the main floor of the Student Services Building. Their telephone number is 515-294-7220 or email accessibility@iastate.edu. Retroactive requests for accommodations will not be honored.

Religious Accommodation: If an academic or work requirement conflicts with your religious practices and/or observances, you may request reasonable accommodations. Your request must be in writing, and your instructor or supervisor will review the request. You or your instructor may also seek assistance from the Dean of Students Office (`http://www.dso.iastate.edu/`) or the Office of Equal Opportunity and Compliance (`http://www.eoc.iastate.edu/`).

Dead Week: This class follows the Iowa State University Dead Week policy as noted in section 10.6.4 of the Faculty Handbook (`http://www.provost.iastate.edu/resources/faculty-handbook`).

Harassment and Discrimination: Iowa State University strives to maintain our campus as a place of work and study for faculty, staff, and students that is free of all forms of prohibited discrimination and harassment based upon race, ethnicity, sex (including sexual assault), pregnancy, color, religion, national origin, physical or mental disability, age, marital status, sexual orientation, gender identity, genetic information, or status as a U.S. veteran. Any student who has concerns about such behavior should contact his/her instructor, Student Assistance (`http://www.dso.iastate.edu/sa/`) at 515-294-1020 or email dso-sas@iastate.edu, or the Office of Equal Opportunity and Compliance (`http://www.eoc.iastate.edu/`) at 515-294-7612.

Accessibility Statement: Iowa State University is committed to assuring that all educational activities are free from discrimination and harassment based on disability status. Students requesting accommodations for a documented disability are required to work directly with staff in Student Accessibility Services (SAS) to establish eligibility and learn about related processes before accommodations will be identified. After eligibility is established, SAS staff will create and issue a Notification Letter for each course listing approved reasonable accommodations. This document will be made available to the student and instructor either electronically or in hard-copy every semester. Students and instructors are encouraged to review contents of the Notification Letters as early in the semester as possible to identify a specific, timely plan to deliver/receive the indicated accommodations. Reasonable accommodations are not retroactive in nature and are not intended to be an unfair advantage. Additional information or assistance is available online at `www.sas.dso.iastate.edu`, by contacting SAS staff by email at accessibility@iastate.edu, or by calling 515-294-7220. Student Accessibility Services is a unit in the Dean of Students Office located at 1076 Student Services Building.

Homework Assignments: There will be four homework assignments. You will have between 10 and 14 days to finish each task. Each homework can be completed using only Microsoft Excel and the virtualized software that will be made available for download during the course. The purpose of this software is to expose preinstalled versions of popular machine learning libraries through a unified user interface.

Each of the four homeworks will require downloading CSV files from the course homepage or other web sites, processing them using standard machine learning algorithms (e.g., using the course-provided virtual machine), analyzing the results in Microsoft Excel. For each homework, there will be an HTML template that will need to be filled out to summarize the results of your experiments.

It is possible to complete each homework without using the course-provided software, e.g., by writing programs that solve problems or by using third-party software packages. In these cases, it is expected that a student describes the solution in a way that enables the course instructor to reproduce it without undue experimentation. For example, if a student writes a program that solves a homework problem, then the source code for this program and a brief description of how to run it need to be included in the solution. The grade for the assignment may be affected by the description quality and by the reproducibility of the results.

There will also be a separate introductory homework called 'Homework 0' that will be made available during the first week of the course. It will count for 1% of your grade. Its purpose is to help students become familiar with the software that will be used during the remaining part of the course. That is, homework 0 will require only installing the software and taking its screenshots to get the credit.

Final Project: The final project for this class should use learning algorithms to help solve a practical research or design problem or focus on evaluating two or more learning algorithms on a data set of your choice. You can join a small group (no more than 3 students) or work alone. Joining a group, however, is highly recommended because you can complement your skills with the skills of other students and divide the writing tasks. **Students are encouraged to select their project topics as soon as possible.** A written project proposal (5–10 pages) will be due on 7/7. The final project report (11–15 pages) will be due on 8/6. Each group will have to present the results of their project during the last two lectures in the course.

You are free to explore public sources of data sets, e.g., the Kaggle platform, data.gov, and other sources, to find suitable data sets. Some of them may be too large to be usable with the course-provided VM. In many cases it will be necessary to perform data conversion and preprocessing to make the data suitable for the VM. For the project, you can use any implementation of learning algorithms, i.e., you are not required to use the VM to obtain your results. Please check the data set license before making the final selection.

Collaboration Policy: You are encouraged to form groups and discuss the readings for this class. These groups can be larger than the project groups. You can discuss the homework assignments with other students who take this class. However, each student is expected to solve the homework problems on their own. Sharing of solutions, including, but not limited to, the Excel spreadsheets that analyze the results of learning algorithms, is not permitted.

IMPORTANT: Cheating, plagiarism, and other academic misconducts will not be tolerated and will be handled according to the ISU's academic dishonesty procedures, which are posted here:
`http://catalog.iastate.edu/academic_conduct/#academicdishonestytext`

Class Participation: The expectation is that you attend or view the offline recording of each lecture. If you forget to view a lecture recording, then you may miss what happened, including any announcements.

Asking questions during the live lecture is one way to satisfy the course participation requirement. Another way is to submit discussion posts related to the course topics. There will be six types of discussions: 1) "Introductions", 2) "Learning algorithm or dataset of the day", 3) "Homework assignments and suggestions for new problems", 4) "Possible topics for course projects", 5) "Help Forum", and 6) "Other discussions related to learning algorithms". Your goal is to submit at least one post or an extended response in at least one of these categories every week during the course except week 9 (project proposal presentations) and week 13 (final project presentations). **Multiple posts generated during the same week don't carry over to the next week.** They do generate **brownie points**, however.

Only one post in the "Introductions" board can count toward the participation score. A post counts toward the participation score only if it is not a restatement of already-posted information, is on-topic, and is not too short. A single paragraph that consists of 3–4 sentences or more that contain new information and do not simply state an opinion or express emotions is sufficient to satisfy the class participation requirement for **one week**. Please don't copy text from other sources to generate your posts, because the Academic Dishonesty policy applies to them similarly to all other assignments in this course. The discussion boards will be moderated by the instructor.

Grading: Your grade will be calculated as follows:

Homework Assignments:	56% (4 $\times$ 14% each)
Final Project Proposal:	10%
Final Project:	18%
Project proposal presentation:	5%
Final project presentation:	5%
Class participation:	5%
Homework 0:	1%

Tentative Schedule

Week	Day/Date		Topic	Assignment
1	Tuesday	1/26	Introduction to HCI 560X	Homework 0 out
	Thursday	1/28	Technology Demo and K-Means Clustering	
2	Tuesday	2/2	Probabilities, Gaussians, and Central Limit Theorem	Homework 0 due
	Thursday	2/4	Gaussian Mixtures	Homework 1 out
3	Tuesday	2/9	Supervised Learning	
	Thursday	2/11	Linear Support Vector Machines	
4	Tuesday	2/16	Kernel Trick	
	Thursday	2/18	TBD	
5	Tuesday	2/23	Dimensionality Reduction and PCA	Homework 1 due
	Thursday	2/25	Project Ideas Presentations	Homework 2 out
6	Tuesday	3/2	Eigenfaces	
	Thursday	3/4	Nearest Neighbors	
7	Tuesday	3/9	Hidden Markov Models (part 1)	
	Thursday	3/11	Hidden Markov Models (part 2)	
8	Tuesday	3/16	Hidden Markov Models (part 3)	Homework 2 due
	Thursday	3/18	TBD	Homework 3 out
9	Tuesday	3/23	Project Proposal Presentations	
	Thursday	3/25	Project Proposal Presentations	
10	Tuesday	3/30	Neurons, Perceptrons, and Back-Propagation	
	Thursday	4/1	TBD	Homework 3 due
11	Tuesday	4/6	Perceptrons and Back-Propagation	Homework 4 out
	Thursday	4/8	Deep Learning and Convolutional Neural Networks	
12	Tuesday	4/13	Markov Decision Processes	
	Thursday	4/15	Policy Iteration and TD Algorithms	
13	Tuesday	4/20	Kalman Filter and Particle Filters	
	Thursday	4/22	Robotic Manipulation	Homework 4 due
14	Tuesday	4/27	Final Project Presentations	
	Thursday	4/29	Final Project Presentations	Final Report due

Tentative List of Lecture Readings Outside Textbook

Subject	**Readings**
Introduction	Turing, A. Computing machinery and intelligence. *Mind*, **59**(236), 433–460 (1950). McCarthy, J., Minsky, M., Rochester, N., and Shannon, C. A proposal for the Dartmouth summer research project on Artificial Intelligence. August 31, 1955. *AI Magazine*, **27**(4), 12–14 (2006). Norvig, P. Teach yourself programming in ten years. `http://norvig.com/21-days.html` (2001).
K-Means Clustering	Lloyd, S. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, **28**(2), 129–137 (1982).
Gaussian Mixtures	Dempster, A., Laird, N., and Rubin, D. Maximum likelihood from incomplete data via the EM algorithms. *Journal of the Royal Statistical Society: Series B (Methodological)* **39**(1), 1–22 (1977). Alpaydin, E. Soft vector quantization and the EM algorithm. *Neural Networks*, **10**, 467–477 (1998).
Supervised Learning	Ghahramani, Z. Probabilistic machine learning and artificial intelligence. *Nature*, **521**, 452–459 (2015). Valiant, L. A theory of the learnable. *Communications of the ACM*, **27**(11), 1134–1142 (1984).
Support Vector Machines	Burges, C. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, **2**, 121–167 (1998).
Dimensionality Reduction	Landauer, T., Foltz, P., and Laham, D. An introduction to latent semantic analysis. *Discourse Processes*, **25**(2&3), 259–284 (1998). Van der Maaten, L. and Hinton, G. Visualizing data using t-SNE. *Journal of Machine Learning Research*, **9**, 2579–2605 (2008).
Eigenfaces	Turk, M. and Pentland, A. Face recognition using eigenfaces. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 586–591 (1991). Ballard, D. *An Introduction to Natural Computation.* MIT Press. (1997). Chapter 4, pages 71–93.
Nearest Neighbors	Cover, T and Hart, P. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, **13**(1), 21–27 (1967).
Performance Analysis	Demšar, J. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, **7**, 1–30 (2006).

Hidden Markov Models	Ferguson, J. Hidden Markov Analysis: An Introduction. *Symposium on the Application of Hidden Markov Models to Text and Speech.* Chapter 2, 8–15 (1980). Rabiner, L. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, **77**(2), 257–286 (1989).
Backpropagation Algorithm	Rumelhart, D., Hinton, G., and Williams, R. Learning representations by back-propagating errors. *Nature*, **323**, 533–536 (1989). Crick, F. The recent excitement about neural networks. *Nature*, **337**, 129–132 (1989). Lillicrap, T., Santoro, A., Marris, L., and Hinton, G. Backpropagation and the brain. *Nature Reviews Neuroscience*, **21**, 335-346 (2020).
Deep Learning	Hinton, G. and Salakhutdinov, R. Reducing the dimensionality of data with neural networks. *Science*, **313**, 504–507 (2006).
ANN Models	LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W., and Jackel, L. Backpropagation applied to handwritten zip-code recognition. *Neural Computation*, **1**, 541-551 (1989). Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, **9**(8), 1735–1780 (1997).
Reinforcement Learning	Sutton, R. Learning to predict by the methods of temporal differences. *Machine Learning*, **3**, 9–44 (1988). Tesauro, G. Temporal difference learning and TD-Gammon. *Communications of the ACM*, **38**(3), 58–68 (1995).