1. **Course number and name**
   CptS/EE 437: Introduction to Machine Learning

2. **Credits and contact hours**
   3 credits, 3 lecture hours

3. **Instructor's or course coordinator's name**
   Diane Cook

4. **Textbook, title, author, and year**
   H. Daume. 2017. *A Course in Machine Learning*. <http://ciml.info/>.
   Other supplemental materials
   T. Mitchell. *Decision tree learning*. <http://www.cs.princeton.edu/courses/archive/spr07/cos424/ papers/mitchell-dectrees.pdf>.
   A. Ng. Advice for applying machine learning. <http://cs229.stanford.edu/materials/ML-advice.pdf>.
   A. Ng. Supervised learning. <http://cs229.stanford.edu/notes/cs229-notes1.pdf> (Part I, Section 1; Part II Section 5).
   L.I. Smith. A tutorial on principal components analysis. <http://www.cs.otago.ac.nz/cosc453/student_tutorials/principal_components.pdf>
   R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*. <http://incompleteideas.net/book/bookdraft2017nov5.pdf>.

5. **Specific course information**
   a. *Catalog description:* Topics in machine learning including linear models for regression and classification, generative models, support vector machines and kernel methods, neural networks and deep learning, decision trees, unsupervised learning, and dimension reduction.
   b. *Prerequisites or corequisites:* CptS 215, CptS 223, or CptS 233, certified major or minor in Computer Science, Computing Engineering, Electrical Engineering, Software Engineering, or Data Analytics.

6. **Specific goals for the course**
   Following completion of this course, students will
   - Have an understanding of major supervised, unsupervised and reinforcement learning techniques (ABET 3a),
   - Have a basic understanding of evaluation methodologies (ABET 3b, 3c),
   - Have a working knowledge of how to apply machine learning technologies to real-world datasets (ABET 3d, 3e, 3f, 3g),
   - Function effectively as part of a team (5b, 5c, 5d, 5e, 5f)

7. **Brief list of topics to be covered**
   - Supervised learning
   - Unsupervised learning

- Reinforcement learning
- Decision trees
- Limits of learning and inductive bias
- Nearest neighbors
- K-means clustering
- Perceptron, neural networks
- Deep networks, tensor flow, generative adversarial networks
- Decision boundaries, margins
- Evaluation of model performance
- Multi-class classification
- Linear regression, logistic regression
- Loss functions, bias, regularization
- Reinforcement learning