Final Examination

CS 205A: Mathematical Methods for Robotics, Vision, and Graphics (Fall 2013), Stanford University

- The exam runs for 3 hours.
- The exam contains eight problems. You must complete the first problem and six of problems 2-8. **CIRCLE THE PROBLEMS YOU WANT GRADED ON THE CHART BELOW; OTHERWISE WE WILL GRADE THE FIRST SIX QUESTIONS ON WHICH YOU HAVE PROVIDED ANY WRITTEN ANSWER.**
- The exam is closed-book. You may use two double-sided $8^1/_2'' \times 11''$ sheets of notes.
- Write your solutions in the space provided. If you need more space, write on the back of the sheet containing the problem and indicate that you have done so.
- Do not spend too much time on any problem. Read them all before beginning.
- Show your work, as partial credit will be awarded.

Circle the six additional problems you want graded.

Problem	1	2	3	4	5	6	7	8	EC	Total
Score										

The Stanford Honor Code

1. The Honor Code is an undertaking of the students, individually and collectively:

    (a) that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;

    (b) that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.

2. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.

3. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

Signature

Name

SUID

YOU MUST COMPLETE THIS PROBLEM.

Problem 1 (Short answer).

(a) Do all square matrices $A \in \mathbb{R}^{n\times n}$ admit a factorization $A = LU$? Explain. [2 points]

(b) Recall Heun's method for time-stepping ODEs:

$$\vec{y}_{k+1} = \vec{y}_k + \frac{h}{2}(F[\vec{y}_k] + F[\vec{y}_k + hF[\vec{y}_k]]).$$

For the one-variable model ODE $y' = ay$ with $a < 0$, calculate the restriction on h for stability of Heun's method. [2 points]
Note: Justin got this bound slightly wrong in lecture, so double-check your work!

(c) Suppose $A, B \in \mathbb{R}^{n \times n}$ and $\vec{a}, \vec{b} \in \mathbb{R}^n$. Find a linear system of equations satisfied by minima of the energy $\|A\vec{x} - \vec{a}\|_2^2 + \|B\vec{x} - \vec{b}\|_2^2$ with respect to $\vec{x}$. [2 points]

(d) Propose a method for finding the least-norm projection of a vector $\vec{v}$ onto the column space of $A \in \mathbb{R}^{m \times n}$ with $m > n$. [2 points]

(e) Recall the Rayleigh quotient iteration strategy for finding eigenvectors of a matrix:

$$\begin{aligned} \sigma_k &= \frac{\vec{v}_{k-1}^\top A \vec{v}_{k-1}}{\|\vec{v}_{k-1}\|_2^2} \\ \vec{w}_k &= (A - \sigma_k I)^{-1} \vec{v}_{k-1} \\ \vec{v}_k &= \frac{\vec{w}_k}{\|\vec{w}_k\|} \end{aligned}$$

We showed that this strategy can converge much more quickly than the basic power method. Why, however, might it still be more efficient to use the power method in some cases? [2 points]

Problem 2 (Optimization for deconvolution). Suppose we take a grayscale photograph of size $n \times m$ and represent it as a vector $\vec{v} \in \mathbb{R}^{nm}$ of values in $[0, 1]$. We used the wrong lens, however, and our photo is blurry! We wish to use *deconvolution* machinery to undo this effect.

(i) Find the KKT conditions for the following optimization problem: [5 points]

$$\begin{array}{rl} \text{minimize}_{\vec{x} \in \mathbb{R}^{nm}} & \|A\vec{x} - \vec{b}\|_2^2 \\ \text{such that} & 0 \leq x_i \leq 1 \; \forall i \in \{1, \ldots, nm\} \end{array}$$

(ii) Suppose we are given a matrix $G \in \mathbb{R}^{nm \times nm}$ taking sharp images to blurry ones. Propose an optimization in the form of (i) for recovering a sharp image from our blurry $\vec{v}$. [2 points]

(iii) We do not know the operator G, making the model in (ii) difficult to use. Suppose, however, that for each $r \geq 0$ we can write a matrix $G_r \in \mathbb{R}^{nm \times nm}$ approximating a blur with radius r. Using the same camera, we now take k *pairs* of photos $(\vec{v}_1, \vec{w}_1), \ldots, (\vec{v}_k, \vec{w}_k)$, where $\vec{v}_i$ and $\vec{w}_i$ are of the same scene but $\vec{v}_i$ is blurry (taken using the same lens as our original bad photo) and $\vec{w}_i$ is sharp. Propose a nonlinear optimization for approximating r using this data. [3 points]

CIRCLE WHICH SIX OF PROBLEMS 2-8 YOU WANT GRADED ON PAGE 1.

Problem 3 (Quantum mechanics). Suppose we wish to write simulation software for quantum physics. The Schrödinger equation and others involve *complex* numbers in $\mathbb{C}$, so we must extend the machinery we have developed in CS 205A to this case. Recall that a complex number $x \in \mathbb{C}$ can be written as $x = a + bi$, where $a, b \in \mathbb{R}$ and $i = \sqrt{-1}$.

(a) Suppose we wish to solve $A\vec{x} = \vec{b}$, but now $A \in \mathbb{C}^{n\times n}$ and $\vec{x}, \vec{b} \in \mathbb{C}^n$. Explain how a linear solver that takes only *real-valued* systems can be used to solve this equation. [5 points]
Hint: Write $A = A_1 + A_2 i$, where $A_1, A_2 \in \mathbb{R}^{n\times n}$. Similarly decompose $\vec{x}$ and $\vec{b}$. In the end you will solve a $2n \times 2n$ real-valued system.

(b) Suppose we discretize Schrödinger's equation for a particular quantum simulation yielding an ODE $\vec{x}' = A\vec{x}$, for $\vec{x}(t) \in \mathbb{C}^n$ and $A \in \mathbb{C}^{n\times n}$. Furthermore, suppose that A is *self-adjoint* and *negative definite*, that is, A satisfies the following properties:

- Self-adjoint: $a_{ij} = \bar{a}_{ji}$, where $\overline{a + bi} = a - bi$.
- Negative definite: $\bar{\vec{x}}^\top A\vec{x} \leq 0$ (and is real) for all $\vec{x} \in \mathbb{C}^n \backslash \{\vec{0}\}$. Here we define $(\bar{\vec{x}})_i \equiv \bar{x}_i$.

Derive a backward Euler strategy for solving this ODE and show that each step can be carried out using conjugate gradients. [5 points]
Hint: Before discretizing, convert the ODE to a real-valued system by applying the same decomposition as suggested for (a) to $\vec{x}(t)$ and A.

Problem 4 (Polar decomposition). In this problem we will add one more matrix factorization to our linear algebra toolbox and derive an algorithm by N. Higham for its computation. The decomposition is used in animation applications interpolating between motions of a rigid object while projecting out undesirable shearing artifacts.

(a) Show that any matrix $A \in \mathbb{R}^{n\times n}$ can be factored $A = WP$, where W is orthogonal and P is symmetric and positive semidefinite. This factorization is known as the polar decomposition. [2 points]
Hint: Write $A = U\Sigma V^\top$ *and show* $V\Sigma V^\top$ *is positive semidefinite.*

(b) The polar decomposition of an invertible $A \in \mathbb{R}^{n\times n}$ can be computed using a simple iterative scheme:

$$X_0 \equiv A \qquad\qquad X_{k+1} = \frac{1}{2}(X_k + (X_k^{-1})^\top)$$

(i) Use the SVD to write $A = U\Sigma V^\top$, and define $D_k = U^\top X_k V$. Show $D_0 = \Sigma$ and $D_{k+1} = \frac{1}{2}(D_k + (D_k^{-1})^\top)$. [2 points]

(ii) From (i), each D_k is diagonal. If d_{ki} is the i-th diagonal element of D_k, show [2 points]

$$d_{(k+1)i} = \frac{1}{2}\left(d_{ki} + \frac{1}{d_{ki}}\right).$$

(iii) Assume $d_{ki} \to c_i$ as $k \to \infty$ (this convergence assumption requires proof in real life!). Show $c_i = 1$. [2 points]

(iv) Use (iii) to show $X_k \to UV^\top$. [2 points]

Problem 5 (Conjugate gradients).

(a) If we use infinite-precision arithmetic (so rounding is not an issue), can the conjugate gradients algorithm be used to recover *exact* solutions to $A\vec{x} = \vec{b}$ for symmetric positive definite matrices A? Why or why not? [3 points]

(b) Suppose $A \in \mathbb{R}^{n\times n}$ is invertible but not symmetric or positive definite.

(i) Show that $A^\top A$ is symmetric and positive definite. [1 point]

(ii) Propose a strategy for solving $A\vec{x} = \vec{b}$ using the conjugate gradients algorithm based on your observation in (i). [3 points]

(iii) How quickly do you expect conjugate gradients to converge in this case? Why? [3 points]

CIRCLE WHICH SIX OF PROBLEMS 2-8 YOU WANT GRADED ON PAGE 1.

Problem 6 (Interpolation).

(a) List one advantage and one disadvantage of each of the following bases for polynomials of degree n:

(i) Monomial basis [1 point]:

(ii) Lagrange basis [1 point]:

(iii) Newton basis [1 point]:

(b) Suppose we wish to interpolate a function $f : \mathbb{R} \to \mathbb{R}$ given k data points $(x_i, f(x_i))$. When might it be preferable to use piecewise polynomial interpolation over a single polynomial of degree k? What is a possible drawback? [4 points]

(c) Write the degree-four polynomial interpolating between the data points $(-2, 15)$, $(0, -1)$, $(1, 0)$, and $(3, -2)$. [3 points]
Hint: Your answer does not have to be written in the monomial basis.

CIRCLE WHICH SIX OF PROBLEMS 2-8 YOU WANT GRADED ON PAGE 1.

Problem 7 (Quadrature).

(a) Derive α, β, and x_1 such that the following quadrature rule holds exactly for polynomials of degree ≤ 2 : [5 points]

$$\int_0^2 f(x)\,dx \approx \alpha f(0) + \beta f(x_1)$$

(b) What is a composite quadrature rule? When are composite quadrature rules useful? [2 points]

(c) Suppose we are given a quadrature rule of the form $\int_0^1 f(x)\,dx \approx af(0) + bf(1)$ for some $a, b \in \mathbb{R}$. Propose a composite rule for approximating $\int_0^1 f(x)\,dx$ given $n+1$ closed sample points $y_0 \equiv f(0), y_1 \equiv f(1/n), y_2 \equiv f(2/n), \ldots, y_n \equiv f(1)$. [3 points]

Problem 8 (ODE).

(a) Suppose we wish to simulate a spring by solving the ODE $d^2y/dt^2 = -y$ with $y(0) = 0$ and $y'(0) = 1$. We obtain the three plots of $y(t)$ below by using forward Euler, backward Euler, and symplectic Euler. Label which plot is which. [3 points]

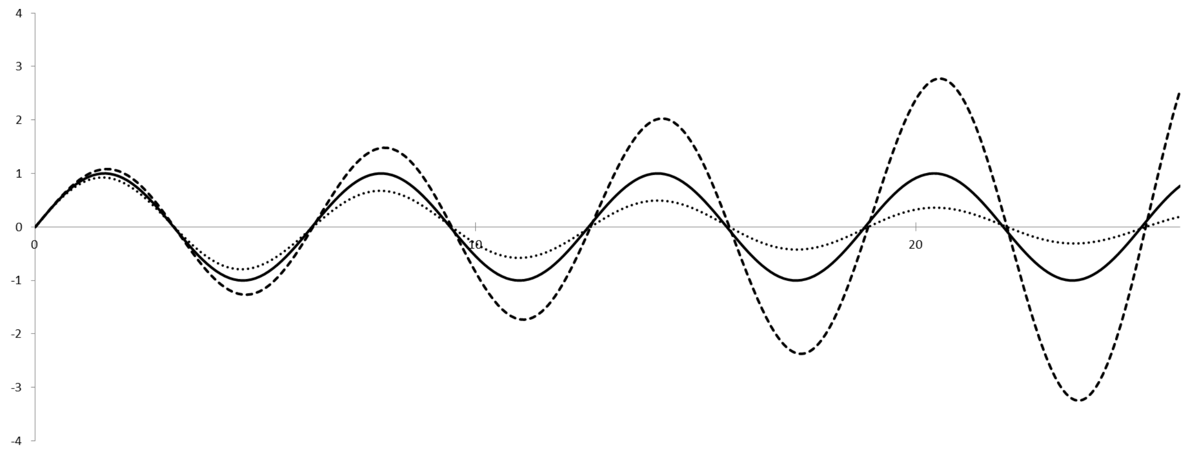

(b) Suppose we wish to solve the ODE $dy/dt = -\sin y$ numerically. For time step $h > 0$, write the implicit backward Euler equation for approximating y_{k+1} at $t = (k+1)h$ given y_k at $t = kh$. [3 points]

(c) Write the Newton iteration for solving the equation from part (b) for y_{k+1}. [4 points]

THIS EXTRA CREDIT PROBLEM IS COMPLETELY OPTIONAL.

Extra credit. A few open-ended conceptual problems on partial differential equations.

(a) Suppose we wish to solve the Laplace equation $\nabla^2 f = 0$ on the unit disc $D = \{(x, y) : x^2 + y^2 \leq 1\}$. We will impose Dirichlet boundary conditions $f(x, y) \equiv g(x, y)$ for all (x, y) with $x^2 + y^2 = 1$, where $g(x, y)$ is given as input.

The two-dimensional Laplacian operator we introduced in lecture was based on a grid structure:

$$(\nabla^2 y)_{ij} \approx \frac{1}{h}(y_{(i-1)j} + y_{i(j-1)} + y_{(i+1)j} + y_{i(j+1)} - 4y_{ij})$$

Obviously this grid does not align with the boundary of the disc ∂D!

Propose a way to discretize and solve this equation on D by changing our discretization of $\nabla^2 y$ and/or working out reasonable boundary conditions on the grid. [5 points]

(b) In CS 205A we studied solutions of *differential* equations, but another more advanced area of analysis involves *integral* equations. A "Fredholm equation of the first type" is given by:

$$f(x) = \int_a^b K(x, t)\phi(t)\, dt,$$

where $K : [a, b] \times [a, b] \to \mathbb{R}$ and $f : [a, b] \to \mathbb{R}$ are given and $\phi : [a, b] \to \mathbb{R}$ is unknown. Propose a discretization for approximating ϕ. Is your matrix sparse? [5 points]