The Mathematics Institutes' MODERN MATH WORKSHOP

How Can Data and Models Work Together in Forecasting?

JUAN M. RESTREPO *Group Leader Uncertainty Quantification Group (UQG)*

Mathematics Department
Atmospheric Sciences and Physics Departments
University of Arizona

September, 2010

How do we make predictions in the face of uncertainty?

- We have a model for the dynamics, but it might have inherent errors...
- We have measurements, but these are not complete and there might be measurement errors...

Three Time-dependent Estimation Problems

Given a random time series $\{X(t) \in \mathbb{R}^N : t \leq t_0\}$ (from models, data, controls):

- Retrodiction:

$$\tilde{X}(t): \quad t \leq t_0.$$

  e.g., paleoclimate reconstruction, optimal control path.

- Nudiction:

$$\tilde{X}(t): \quad t = t_0.$$

  e.g., best initial conditions for weather prediction, optimal configuration.

- Prediction:

$$\tilde{X}(t): \quad t > t_0.$$

  e.g., weather prediction, system forecast.

Automating Navigation: flying airplanes and spacecraft, driving rovers and probes...

Image, courtesy of JPL, Pasadena CA.

The Prediction Problem (Methodology/unconstrained data)

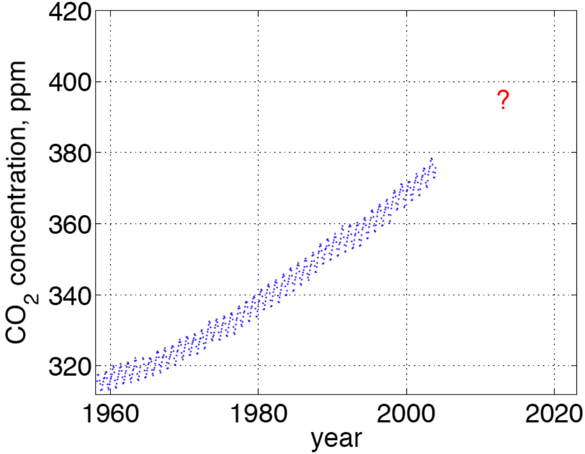

Atmospheric CO_2 at Mauna Loa Observatory (collected by D. Keeling, Scripps).

The Prediction Problem (Methodology/unconstrained data)

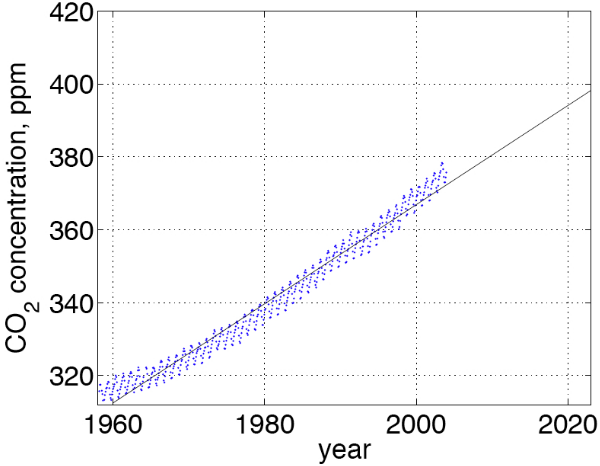

The Prediction Problem (Methodology/unconstrained data)

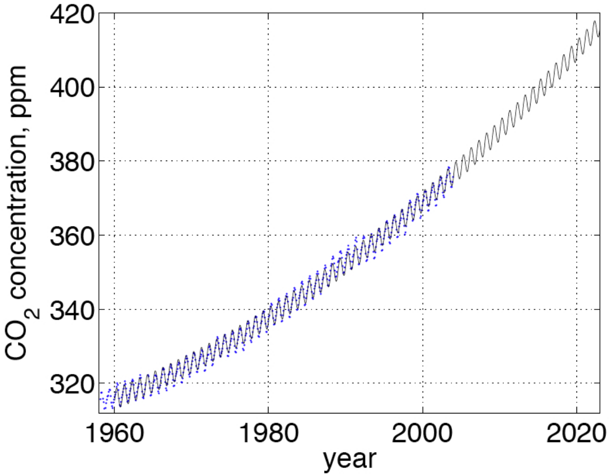

The Prediction Problem (Methodology/unconstrained data)

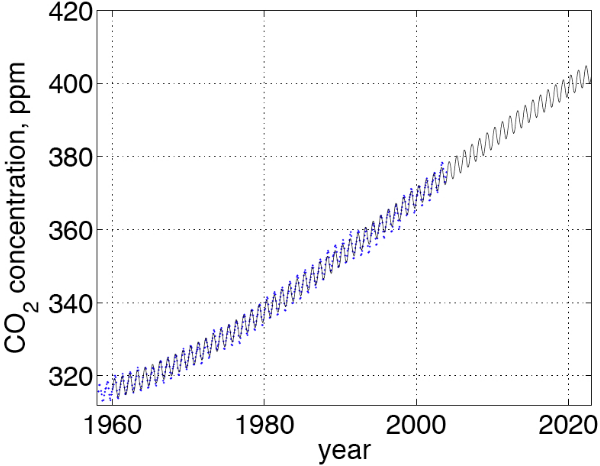

The Prediction Problem

When data fool us...

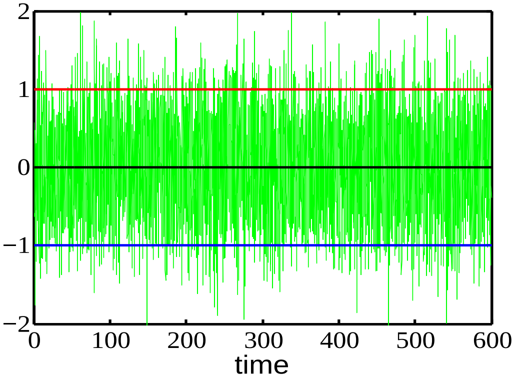

The Prediction Problem

When data fool us...

same data, zoomed in

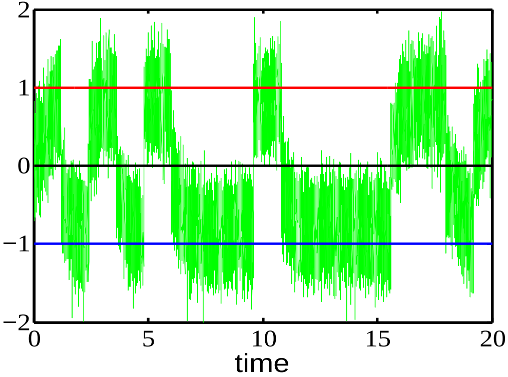

The Prediction Problem

...use our understanding of the dynamics

$$\begin{aligned} dx &= 4x(1-x^2)dt + \kappa dW_t \\ x(0) &= x_0 \end{aligned}$$

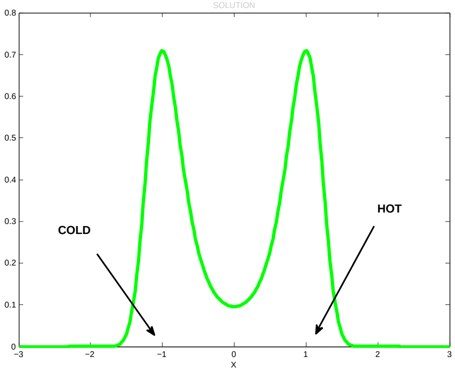

PART I: LINEAR ALGEBRA BACKGROUND

Introduction largely drawn from G. Strang's *Linear Algebra and its Applications* book.

1. Matrices and vectors

- An $m \times n$ matrix is an array with m rows and n columns. It is typically written in the form

$$A = [a_{ij}] = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix},$$

  where i is the row index and j is the column index.

- A column vector is an $m \times 1$ matrix. Similarly, a row vector is a $1 \times n$ matrix.

- The entries a_{ij} of a matrix A may be real or complex.

Matrices and vectors (continued)

- **Examples:**
  - $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ is a 2×2 square matrix with real entries.
  - $u = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$ is a column vector of A.
  - $B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & i & 0 \\ 0 & 0 & 3-7i \end{bmatrix}$ is a 3×3 diagonal matrix, with complex entries.
  - An $n \times n$ diagonal matrix whose entries are all ones is called the $n \times n$ identity matrix.
  - $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$ is a 2×4 matrix with real entries.

Matrix addition and scalar multiplication

Let $A = [a_{ij}]$ and $B = [b_{ij}]$ be two $m \times n$ matrices, and let c be a scalar.

- The matrices A and B are equal if and only if they have the same entries,

$$A = B \Longleftrightarrow a_{ij} = b_{ij}, \text{ for all } i, j, \ 1 \leq i \leq m, \ 1 \leq j \leq n.$$

- The sum of A and B is the $m \times n$ matrix obtained by adding the entries of A to those of B,

$$A + B = [a_{ij} + b_{ij}].$$

- The product of A with the scalar c is the $m \times n$ matrix obtained by multiplying the entries of A by c,

$$cA = [c\, a_{ij}].$$

2. Matrix multiplication

- Let $A = [a_{ij}]$ be an $m \times n$ matrix and $B = [b_{ij}]$ be an $n \times p$ matrix. The product $C = AB$ of A and B is an $m \times p$ matrix whose entries are obtained by multiplying each row of A with each column of B as follows:

$$c_{ij} = \sum_{k=1}^{n} a_{ik} b_{kj}.$$

- **Examples:** Let $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ and $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$.
  - Is the product AC defined? If so, evaluate it.
  - Same question with the product CA.
  - What is the product of A with the third column vector of C?

Matrix multiplication (continued)

- **More examples:**
  - Consider the system of equations

$$\begin{cases} 3x_1 + 2x_2 - x_3 = 4 \\ x_2 - 7x_3 = 0 \\ -x_1 + 4x_2 - 6x_3 = -10 \end{cases}.$$

    Write this system in the form $AX = Y$, where A is a matrix and X and Y are two column vectors.

  - Let

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix}.$$

    Calculate the products AB and BA.

3. Rules for matrix addition and multiplication

- The rules for matrix addition and multiplication by a scalar are the same as the rules for addition and multiplication of real or complex numbers.
- In particular, if A and B are matrices and c_1 and c_2 are scalars, then

$$
\begin{aligned}
A+B &= B+A \\
(A+B)+C &= A+(B+C) \\
c_1\,(A+B) &= c_1\,A+c_1\,B \\
(c_1+c_2)A &= c_1A+c_2A \\
c_1\,(c_2\,A) &= (c_1\,c_2)A
\end{aligned}
$$

whenever the above quantities make sense.

Rules for matrix addition and multiplication (continued)

- The product of two matrices is associative and distributive, i.e.

$$A(BC) = (AB)C = ABC$$
$$A(B+C) = AB + AC \qquad (A+B)C = AC + BC.$$

- However, the product of two matrices is not commutative. If A and B are two square matrices, we typically have

$$AB \neq BA$$

- For two square matrices A and B, the commutator of A and B is defined as

$$[A,B] = AB - BA.$$

In general, $[A,B] \neq 0$. If $[A,B] = 0$, one says that the matrices A and B commute.

4. Transposition

- The transpose of an $m \times n$ matrix A is the $n \times m$ matrix A^T obtained from A by switching its rows and columns, i.e.

$$\text{if} \quad A = [a_{ij}], \qquad \text{then} \quad A^T = [a_{ji}].$$

- **Example:** Find the transpose of $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$.
- Some properties of transposition. If A and B are matrices, and c is a scalar, then

$$(A+B)^T = A^T + B^T \qquad (cA)^T = cA^T$$
$$(AB)^T = B^T A^T \qquad \left(A^T\right)^T = A,$$

whenever the above quantities make sense.

Linear independence

- A linear combination of the n vectors $a_1,\ a_2,\ \cdots,\ a_n$ is an expression of the form
$$c_1a_1 + c_2a_2 + \cdots + c_na_n,$$
where the c_i's are scalars.

- A set of vectors $\{a_1, a_2, \cdots, a_n\}$ is linearly independent if the only way of having a linear combination of these vectors equal to zero is by choosing all of the coefficients equal to zero. In other words, $\{a_1, a_2, \cdots, a_n\}$ is linearly independent if and only if
$$c_1a_1 + c_2a_2 + \cdots + c_na_n = 0 \Longrightarrow c_1 = c_2 = \cdots = c_n = 0.$$

Linear independence (continued)

- **Examples:**
  - Are the columns of the matrix $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ linearly independent?
  - Same question with the columns of the matrix $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$.
  - Same question with the rows of the matrix C defined above.
- A set that is not linearly independent is called linearly dependent.
- Can you find a condition on a set of n vectors, which would guarantee that these vectors are linearly dependent?

6. Vector space

- A real (or complex) vector space is a non-empty set V whose elements are called vectors, and which is equipped with two operations called vector addition and multiplication by a scalar.

- The vector addition satisfies the following properties.

  The sum of two vectors $a \in V$ and $b \in V$ is denoted by $a+b$ and is an element of V.

  It is commutative: $a+b=b+a$, for all $a,b \in V$.

  It is associative: $(a+b)+c=a+(b+c)$ for all $a,b,c \in V$.

  There exists a unique zero vector, denoted by 0, such that for every vector $a \in V$, $a+0=a$.

  For each $a \in V$, there exists a unique vector $(-a) \in V$ such that $a+(-a)=0$.

Vector space (continued)

- The multiplication by a scalar satisfies the following properties.

  The multiplication of a vector $a \in V$ by a scalar $\alpha \in \mathbb{R}$ (or $\alpha \in \mathbb{C}$) is denoted by $\alpha\, a$ and is an element of V.

  Multiplication by a scalar is distributive:

$$\alpha\,(a+b) = \alpha\, a + \alpha\, b, \qquad (\alpha+\beta)\, a = \alpha\, a + \beta\, a,$$

  for all $a, b \in V$ and $\alpha, \beta \in \mathbb{R}$ (or $\mathbb{C}$).

  It is associative: $\alpha\,(\beta a) = (\alpha\,\beta)\, a$ for all $a \in V$ and $\alpha, \beta \in \mathbb{R}$ (or $\mathbb{C}$).

  Multiplying a vector by 1 gives back that vector, i.e.

$$1\, a = a,$$

  for all $a \in V$.

Bases and dimension

- The span of set of vectors $\mathscr{U} = \{a_1, a_2, \cdots, a_n\}$ is the set of all linear combinations of vectors in $\mathscr{U}$. It is denoted by

$$\text{Span}\{a_1, a_2, \cdots, a_n\} \text{ or } \text{Span}(\mathscr{U})$$

  and is a subspace of V.

- A basis $\mathscr{B}$ of a subspace S of V is a set of vectors of S such that
  - $\text{Span}(\mathscr{B}) = S$;
  - $\mathscr{B}$ is a linearly independent set.

- Theorem: If a basis $\mathscr{B}$ of a subspace S of V has n vectors, then all other bases of S have exactly n vectors.

- The dimension of a vector space V (or of a subspace S of V) spanned by a finite number of vectors is the number of vectors in any of its bases.

Rank

- The row space of an $m \times n$ matrix A is the span of the row vectors of A. If A has real entries, the row space of A is a subspace of $\mathbb{R}^n$.
- Similarly, the column space of A is the span of the column vectors of A, and is a subspace of $\mathbb{R}^m$.
- The rank of a matrix A is the dimension of its column space.
- Theorem: The dimensions of the row and column spaces of a matrix A are the same. They are equal to the rank of A.
- **Example:** Check that the row and column spaces of $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$ are vector subspaces, and find their dimension.

The rank theorem

- The null space of an $m \times n$ matrix A, $\mathscr{N}(A)$ is the set of vectors u such that $Au = 0$. If A has real entries, then $\mathscr{N}(A)$ is a subspace of $\mathbb{R}^n$.
- The rank theorem states that if A is an $m \times n$ matrix, then

$$\mathrm{rank}(A) + \dim(\mathscr{N}(A)) = n.$$

- **Example:** Find the rank and the null space of the matrix $C = \begin{bmatrix} 1 & 2 & 3 & 10 \\ 1 & 6 & -8 & 0 \end{bmatrix}$.
  Check that the rank theorem applies.

Linear systems of equations

- A linear system of equations of the form

$$\begin{aligned}
a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\
a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\
\ldots & \\
a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m
\end{aligned}$$

can be written in matrix form as $AX = B$, where

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}, \quad X = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}, \quad B = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}$$

Solution(s) of a linear system of equations

- Given a matrix A and a vector B, a solution of the system $AX = B$ is a vector X which satisfies the equation $AX = B$.
- If B is not in the column space of A, then the system $AX = B$ has no solution. One says that the system is not consistent. In the statements below, we assume that the system $AX = B$ is consistent.
- If the null space of A is non-trivial, then the system $AX = B$ has more than one solution.
- The system $AX = B$ has a unique solution provided $\dim(\mathscr{N}(A)) = 0$.
- Since, by the rank theorem, $\mathrm{rank}(A) + \dim(\mathscr{N}(A)) = n$ (recall that n is the number of columns of A), the system $AX = B$ has a unique solution if and only if $\mathrm{rank}(A) = n$.

Solution(s) of a linear system of equations (continued)

- A linear system of the form $AX = 0$ is said to be homogeneous.
- Solutions of $AX = 0$ are vectors in the null space of A.
- If we know one solution X_0 to $AX = B$, then all solutions to $AX = B$ are of the form
$$X = X_0 + X_h$$
where X_h is a solution to the associated homogeneous equation $AX = 0$.
- In other words, the general solution to the linear system $AX = B$, if it exists, can be written as the sum of a particular solution X_0 to this system, plus the general solution of the associated homogeneous system.

2. Inverse of a matrix

- If A is a square $n \times n$ matrix, its inverse, if it exists, is the matrix, denoted by A^{-1}, such that
$$AA^{-1} = A^{-1}A = I_n,$$
where I_n is the $n \times n$ identity matrix.
- A square matrix A is said to be singular if its inverse does not exist. Similarly, we say that A is non-singular or invertible if A has an inverse.
- The inverse of a square matrix $A = [a_{ij}]$ is given by
$$A^{-1} = \frac{1}{\det(A)} [C_{ij}]^T,$$
where $\det(A)$ is the determinant of A and C_{ij} is the matrix of cofactors of A.

Determinant of a matrix

- The determinant of a square $n \times n$ matrix $A = [a_{ij}]$ is the scalar

$$\det(A) = \sum_{i=1}^{n} a_{ij} C_{ij} = \sum_{j=1}^{n} a_{ij} C_{ij}$$

where the cofactor C_{ij} is given by

$$C_{ij} = (-1)^{i+j} M_{ij},$$

and the minor M_{ij} is the determinant of the matrix obtained from A by "deleting" the i-th row and j-th column of A.

- **Example:** Calculate the determinant of $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$.

Properties of determinants

- If a determinant has a row or a column entirely made of zeros, then the determinant is equal to zero.
- The value of a determinant does not change if one replaces one row (resp. column) by itself plus a linear combination of other rows (resp. columns).
- If one interchanges 2 columns in a determinant, then the value of the determinant is multiplied by -1.
- If one multiplies a row (or a column) by a constant C, then the determinant is multiplied by C.
- If A is a square matrix, then A and A^T have the same determinant.

Properties of the inverse

- Since the inverse of a square matrix A is given by

$$A^{-1} = \frac{1}{\det(A)} \left[C_{ij}\right]^T,$$

  we see that A is invertible if and only if $\det(A) \neq 0$.

- If A is an invertible 2×2 matrix, $\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$, then

$$A^{-1} = \frac{1}{\det(A)} \begin{bmatrix} a_{22} & -a_{12} \\ -a_{21} & a_{11} \end{bmatrix},$$

  and $\det(A) = a_{11}a_{22} - a_{21}a_{12}$.

- If A and B are invertible, then

$$(AB)^{-1} = B^{-1}A^{-1} \qquad \text{and} \qquad \left(A^{-1}\right)^{-1} = A.$$

Linear systems of n equations with n unknowns

- Consider the following linear system of n equations with n unknowns,

$$\begin{aligned} &a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n = b_1 \\ &a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n = b_2 \\ &\ldots \\ &a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nn}x_n = b_n \end{aligned}$$

- This system can be also be written in matrix form as $AX = B$, where A is a square matrix.
- If $\det(A) \neq 0$, then the above system has a unique solution X given by

$$X = A^{-1}B.$$

Linear systems of equations - summary

Consider the linear system $AX = B$ where A is an $m \times n$ matrix.

- The system may not be consistent, in which case it has no solution.
- To decide whether the system is consistent, check that B is in the column space of A.
- If the system is consistent, then
  - Either $\text{rank}(A) = n$ (which also means that $\dim(\mathcal{N}(A)) = 0$), and the system has a unique solution.
  - Or $\text{rank}(A) < n$ (which also means that $\mathcal{N}(A)$ is non-trivial), and the system has an infinite number of solutions.

Linear systems of equations - summary (continued)

Consider the linear system $AX = B$ where A is an $m \times n$ matrix.

- If $m = n$ and the system is consistent, then
  - Either $\det(A) \neq 0$, in which case $\text{rank}(A) = n$, $\dim(\mathcal{N}(A)) = 0$, and the system has a unique solution;
  - Or $\det(A) = 0$, in which case $\dim(\mathcal{N}(A)) > 0$, $\text{rank}(A) < n$, and the system has an infinite number of solutions.
- Note that when $m = n$, having $\det(A) = 0$ means that the columns of A are linearly dependent.
- It also means that $\mathcal{N}(A)$ is non-trivial and that $\text{rank}(A) < n$.

3. Eigenvalues and eigenvectors

- Let A be a square $n \times n$ matrix. We say that X is an eigenvector of A with eigenvalue λ if

$$X \neq 0 \qquad \text{and} \qquad AX = \lambda X.$$

- The above equation can be re-written as

$$(A - \lambda I_n)X = 0.$$

- Since $X \neq 0$, this implies that $A - \lambda I_n$ is not invertible, i.e. that $\det(A - \lambda I_n) = 0$.
- The eigenvalues of A are therefore found by solving the characteristic equation $\det(A - \lambda I_n) = 0$.

Eigenvalues

- The characteristic polynomial $\det(A-\lambda I_n)$ is a polynomial of degree n in λ. It has n complex roots, which are not necessarily distinct from one another.
- If λ is a root of order k of the characteristic polynomial $\det(A-\lambda I_n)$, we say that λ is an eigenvalue of A of algebraic multiplicity k.
- If A has real entries, then its characteristic polynomial has real coefficients. As a consequence, if λ is an eigenvalue of A, so is $\bar{\lambda}$.
- It A is a 2×2 matrix, then its characteristic polynomial is of the form $\lambda^2-\lambda\,\mathrm{Tr}(A)+\det(A)$, where the trace of A, $\mathrm{Tr}(A)$, is the sum of the diagonal entries of A.

Eigenvalues (continued)

- **Examples:** Find the eigenvalues of the following matrices.
  - $A = \begin{bmatrix} -1 & 0 \\ 0 & 5 \end{bmatrix}.$
  - $B = \begin{bmatrix} -1 & 9 \\ 0 & 5 \end{bmatrix}.$
  - $C = \begin{bmatrix} -13 & -36 \\ 6 & 17 \end{bmatrix}.$
  - $D = \begin{bmatrix} 4 & -1 & 1 \\ -1 & 4 & -1 \\ -1 & 1 & 2 \end{bmatrix}.$

Eigenvectors

- Once an eigenvalue λ of A has been found, one can find an associated eigenvector, by solving the linear system

$$(A - \lambda I_n)X = 0.$$

- Since $\mathcal{N}(A - \lambda I_n)$ is not trivial, there is an infinite number of solutions to the above equation. In particular, if X is an eigenvector of A with eigenvalue λ, so is αX, where $\alpha \in \mathbb{R}$ (or $\mathbb{C}$) and $\alpha \neq 0$.
- The set of eigenvectors of A with eigenvalue λ, together with the zero vector, form a subspace of $\mathbb{R}^n$ (or $\mathbb{C}^n$), E_λ, called the eigenspace of A corresponding to the eigenvalue λ.
- The dimension of E_λ is called the geometric multiplicity of λ.

Eigenvectors (continued)

- **Examples:** Find the eigenvectors of the following matrices. Each time, give the algebraic and geometric multiplicities of the corresponding eigenvalues.
  - $A = \begin{bmatrix} -1 & 0 \\ 0 & 5 \end{bmatrix}$.
  - $C = \begin{bmatrix} -13 & -36 \\ 6 & 17 \end{bmatrix}$.
  - $D = \begin{bmatrix} 4 & -1 & 1 \\ -1 & 4 & -1 \\ -1 & 1 & 2 \end{bmatrix}$.

Properties of eigenvalues and eigenvectors

- The geometric multiplicity m_λ of an eigenvalue λ is less than or equal to its algebraic multiplicity M_λ.
- If $M_\lambda = 1$, then $m_\lambda = 1$.
- If m_λ is not equal to M_λ, then one can find $M_\lambda - m_\lambda$ linearly independent generalized eigenvectors of A, by solving a sequence of equations of the form

$$(A - \lambda I_n)\, U_{i+1} = U_i, \qquad i \in \{1, \cdots, M_\lambda - m_\lambda\}$$

where $U_1 = X_\lambda$ is a genuine eigenvector of A with eigenvalue λ.

Properties of eigenvalues and eigenvectors (continued)

- **Examples:** Find the genuine and generalized eigenvectors of the following matrices

  - $M = \begin{bmatrix} 4 & 1 & 0 & 0 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 4 & 1 \\ 0 & 0 & 0 & 4 \end{bmatrix}.$

  - $N = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$

- If A has k distinct eigenvalues and $\mathscr{B}_1, \cdots, \mathscr{B}_k$ are bases of the corresponding generalized eigenspaces, then $\{\mathscr{B}_1, \cdots, \mathscr{B}_k\}$ is a basis of $\mathbb{R}^n$ (or $\mathbb{C}^n$).

Linear Transformations

Suppose the vectors $x_1, x_2, \ldots, x_n$ are a basis for the linear vector space V and $y_1, y_2, \ldots, y_m$ are a basis for the linear vector space W. Then each linear transformation A from V to W is represented by a matrix. The j^{th} column is found by applying A to the j^{th} basis vector; the result, Ax_j is a *linear combination* of the y's and the coefficients in that combination go into column j:

$$Ax_j = a_1 y_1 + a_2 y_2 + \ldots + a_m y_m.$$

Important Linear Transformations

The linear transformation Az transforms z as follows:

from $\mathbb{R}^n$ to $\mathbb{R}^m$, where m can be equal to n. Some of the most important linear transformations are (consider $A \in \mathbb{R}^{2\times 2}$ and $z \in \mathbb{R}^2$:

- Dilation: $A = cI_n$, where c is constant. It stretches (or shrinks) x.
- Rotation: $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. Rotates by some angle 90^o (coordinate rotation) while preserving the size of the vector x.
- Reflection: $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Reflects about the axis $y = x$. Generally, reflects about some axis of symmetry.
- Projection: $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. Takes z in 2-dimensions, to 1-dimension: here it takes a vector z in the plane (x,y) to the nearest point $(x,0)$ on the horizontal axis. Note that neither the dimension or the length of z are preserved.

Important Linear Transformations

More generally (again, consider $A \in \mathbb{R}^{2\times 2}$ and $z \in \mathbb{R}^2$:

- Dilation: $A = cI_n$, where c is constant. It stretches (or shrinks) x.
- Rotation by angle θ:
$$A = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}.$$
- Reflection across line at angle θ:
$$A = \begin{bmatrix} 2\cos^2\theta - 1 & 2\cos\theta\sin\theta \\ 2\cos\theta\sin\theta & 2\sin^2\theta - 1 \end{bmatrix}.$$
- Projection onto line at angle θ:
$$A = \begin{bmatrix} \cos^2\theta & \cos\theta\sin\theta \\ \sin\theta\cos\theta & \sin^2\theta \end{bmatrix}.$$

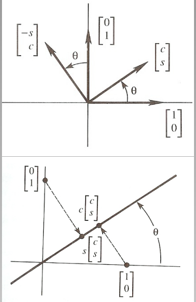

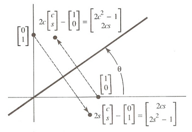

Projections onto a Line

Problem: Given a vector **a** and another vector **b**, the challenge is to find the shortest distance between the tip of one of the vectors to any point colinear with the other vector.

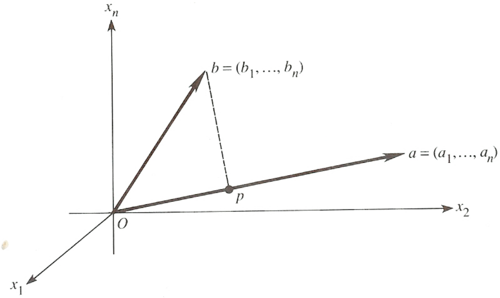

Note that this point p is such that a vector perpendicular to **a** extends to **b**.
This is a first example of least squares problem.

Projections onto a Line

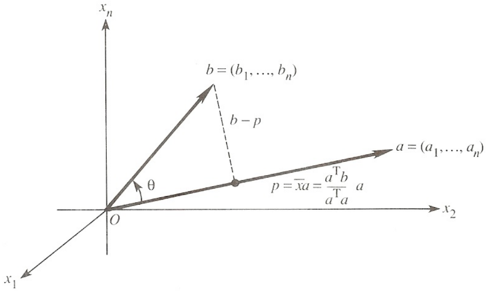

The projection of b into the line, through 0 and a is

$$p = \bar{x}a = \frac{a^\top b}{|a|^2} a.$$

Note that $a^\top b = |a|\,|b| \cos\theta$.

Least Squares in Several Variables

Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^{m}$.

A Practical Problem: Given m observations (data), you want to propose a *model* of the form

$$A\bar{x} - b = r,$$

such that A is as "compact" as possible and/or r is as "small" as possible.

Geometrically, for $m = 3$, and $n = 2$ (thus $\bar{x} \in \mathbb{R}^2$):

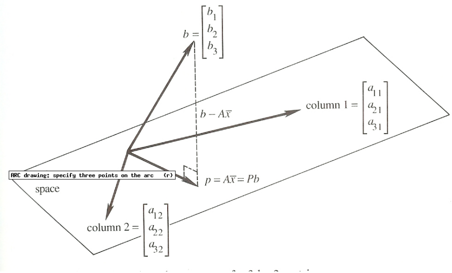

Least Squares in Several Variables

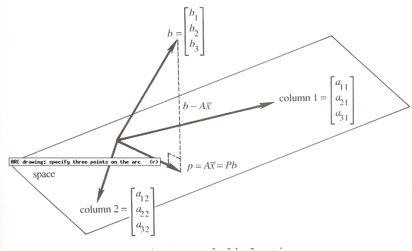

r must be perpendicular to *every* column of A: That is,

$$\begin{aligned} a_1^\top (b - A\overline{x}) &= 0 \\ &\vdots \\ a_n^\top \cdot (b - A\overline{x}) &= 0. \end{aligned}$$

Or

$$A^\top r = 0, \qquad \text{equivalently,} \quad A^\top A\overline{x} = A^\top b.$$

Least Squares in Several Variables

The ”smallness” of r can be measured in terms of a norm.
A convenient norm is the 2-norm:

$$E := ||r||_2^2 = r^\top r = (A\bar{x} - b)^\top (A\bar{x} - b).$$

We note that

$$\frac{1}{2}\frac{dE}{d\bar{x}} = A^\top (A\bar{x} - b)$$

which we call the *normal equations*.

For a given b and a choice of A, we can find $\bar{x}$ which minimizes the distance squared E: E is smallest where $\frac{dE}{d\bar{x}} = 0$. This equation gives us $\bar{x}$.

The Least Squares Solution to the system of m equations in n unknowns

- It satisfies $A^\top A\bar{x} = A^\top b$
- If the columns of A are linearly independent, then $A^\top A$ is invertible and

$$\bar{x} = (A^\top A)^{-1}A^\top b.$$

- The projection of b into the column space of A is thus

$$p = A\bar{x} = A\,(A^\top A)^{-1}A^\top b.$$

Note: $A^\top A$ is symmetric and has the same null space as A, invertible if the columns of A are linearly independent.

Example: Given data: $b = 1$ at $t = -1$, $b = 1$ at $t = 1$, $b = 3$ at $t = 2$. Propose a model of the form $D + Gt_i = b_i$. Find scalars D and G, that in the least-square sense satisfies the equation, for all data points. Solution:

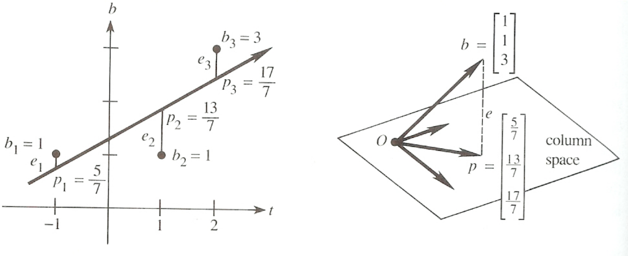

Try next a model of the form $D + Gt_i^2 = b_i$.

The Gaussian Probability Distribution

$$p(X) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{1}{2\sigma^2}(X-m)^2}$$

is a 2-parameter probability distribution:

$$m = \int_{-\infty}^{\infty} x p(x)\, dx := \langle x \rangle,$$

and

$$\sigma = \int_{-\infty}^{\infty} (x-m)^2 p(x)\, dx := \langle (x-m)^2 \rangle.$$

m and σ^2 are known as the *mean* and *variance* (or the first and second moments of $p(X)$).

Gaussian Probability Distributions in Vector Spaces

Let $x := [x_1 \, x_2 \dots x_N]^\top$, where x_i are scalars with Gaussian PDF's. Let $m \in \mathbb{R}^N$

$$p(X) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{1}{2\sigma^2}(X-m)^2}$$

is a 2-parameter vector probability distribution:

$$m = \int_{-\infty}^{\infty} x p(x)\, dx = \langle x \rangle,$$

and

$$C := \int_{-\infty}^{\infty} (x-m)(x-m)^\top p(x)\, dx = \langle (x-m)(x-m)^\top \rangle$$

m and $C \in \mathbb{R}^{+n\times n}$ are known as the *mean* and *variance* (or the first and second moments of $p(X)$). Here,

$$p(X) = \frac{1}{(2\pi)^{N/2}} \frac{1}{\sqrt{\det C}} e^{-\frac{1}{2}[X-m]^\top C^{-1}[X-m]}.$$

Suppose C is diagonal

$$C = \langle x_i x_j \rangle = \delta_{ij} \sigma_i^2,$$

and $m = 0$, then the normal, delta-correlated vector distribution is

$$p(X) = \frac{1}{(2\pi)^{N/2}} \frac{1}{\sqrt{\det C}} e^{-\frac{1}{2}\left[\frac{x_1^2}{\sigma_1^2} + \frac{x_2^2}{\sigma_2^2} + \ldots + \frac{x_N^2}{\sigma_N^2}\right]}$$

This normal distribution is known as *vector white noise*.

Back to Least Squares, a Statistical Interpretation

Consider data $b(t_i) = L(t_i) + n(t_i)$, $i = 1, 2, ..., m$.

$$L(t_i) = D + Gt_i^{\alpha} := A_{ij}X_j,$$

here $X := [D\ G]^\top$. α is a parameter associated with the "model." Succinctly:

$$AX - b = N.$$

Assume that the Gaussian noise processes are have zero mean and are "delta-correlated": $\langle n(t_i) \rangle = 0$ and $\langle n(t_i)n(t_j) \rangle = \delta_{ij}\sigma_i^2$.
So the Least Squares gives an estimate $\tilde{x}$, given by $\tilde{x} = (A^\top A)^{-1}A^\top b$, with error covariance

$$P := \langle (x - \tilde{x})(x - \tilde{x})^\top \rangle = (A^\top A)^{-1}A^\top < NN^\top > A(A^\top A)^{-1} = \sigma_i^2 (A^\top A)^{-1}\delta ij.$$

and estimated fit $\tilde{N} = b - A\tilde{x} = (\delta_{ij} - A(A^\top A)^{-1}A^\top)b$.

When there's excellent data

When data do not fail us...

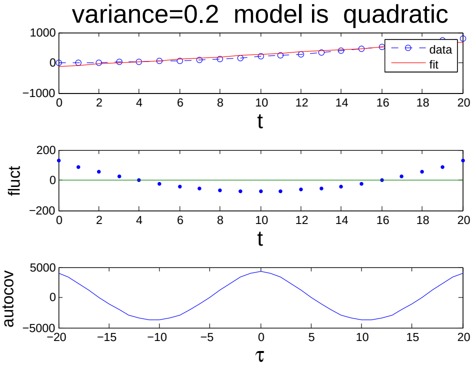

When data do not fail us...

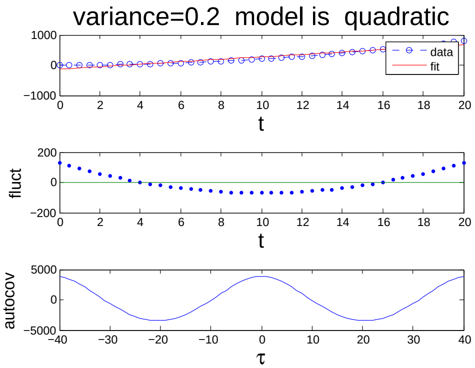

When data do not fail us...

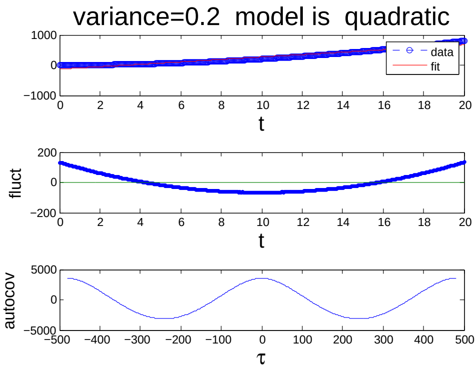

When data do not fail us...

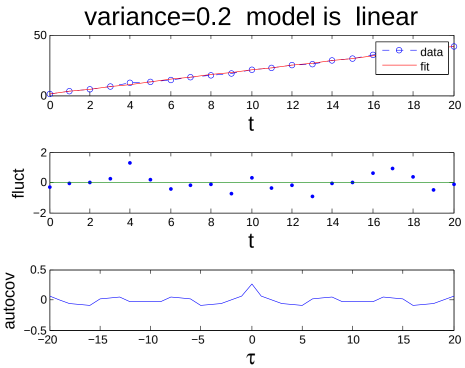

When data do not fail us...

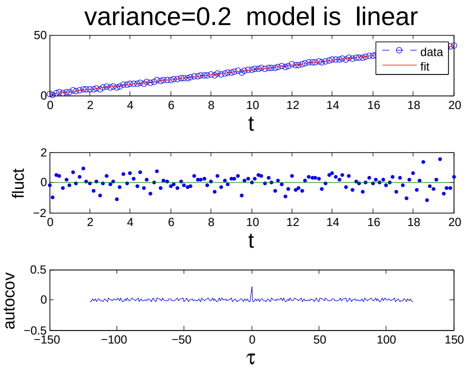

When data do not fail us...

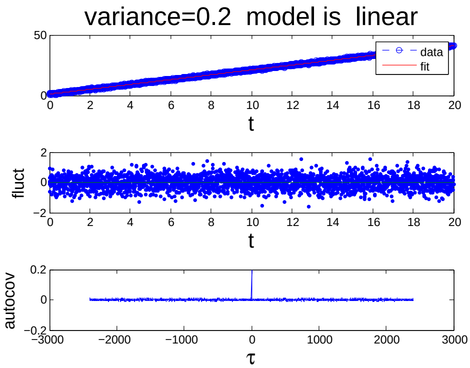

Combining Observations and Mathematical Models

Why is this a good idea? Suppose the data for some experiment was:

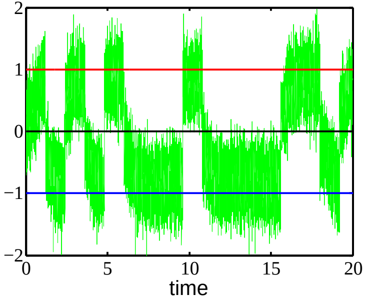

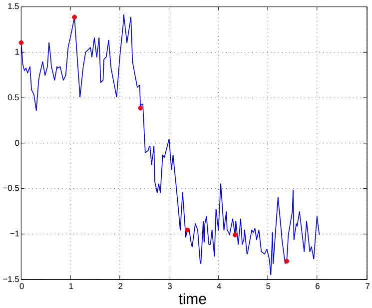

If you used only the model...

$$dx = 4x(1-x^2)dt + \kappa dW_t$$
$$x(0) = x_0$$

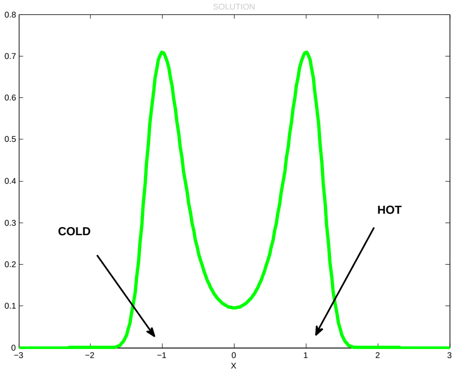

Using Data and Models:

Focus on linear-Gaussian case:

- $P(x|y) \propto \text{Likelihood} \times \text{Prior}.$
- Use data y for likelihood: $y = Hx + n_1$
- Use model for prior: $Ax - b = n_2$

Combining Models and Data, the Linear Gaussian Case

if $n_1 \sim e^{-\xi^2/Q^2}$ and $n_2 \sim e^{-\zeta^2/R^2}$ are normally distributed

$$P(x|y) \sim e^{-\frac{(Ax-b)^2}{Q^2}} e^{-\frac{(y-Hx)^2}{R^2}} = e^{-\left[\frac{(Ax-b)^2}{Q^2} + \frac{(y-Hx)^2}{R^2}\right]}$$

BAYESIAN Least Squares for Linear/Gaussian Problems

Linear/Gaussian global data assimilation: given a model

$$A(m)x - b = \theta_1,$$

and data

$$B(m)x - y = \theta_2$$

Leads to the following least squares problem:

$$W(m)x - V = \Theta,$$
$$\Theta \sim \mathcal{N}(0, R).$$

Find $\tilde{x}$ *, mean, such that* $\mathbb{E}(\theta^\top \theta)$ *is minimized.*
(Also, *find the uncertainty* $P := \mathbb{E}[(x - \tilde{x})(x - \tilde{x})^\top]$).
Remark: Minimizing the variance above, maximizes the Bayesian conditional probability:

$$P(x|y) \propto \exp(-\Theta^2/R) = \exp(-\theta_2^2/r_2)\exp(\theta_1^2/r_1).$$

Recalling Least Squares

Given the least squares problem:

$$Wx - V = \Theta,$$

Extremize

$$J := \|\Theta^\top \Theta\| = \|[Wx - V]^\top [Wx - V]\|.$$

Solve the Normal Equations $W^\top W\tilde{x} = W^\top V$, which yield

$$\begin{aligned} \tilde{x} &= (W^\top W)^{-1} W^\top V, \quad \text{the estimate,} \\ \tilde{n} &= V - W\tilde{x} = (I - W(W^\top W)^{-1} W^\top)V, \text{the residual,} \end{aligned}$$

Same, but ROW WEIGHTED Least Squares

Given the least squares problem:

$$\mathcal{W}x - \mathcal{V} = \boldsymbol{\Theta},$$

with Normal $\langle \boldsymbol{\Theta} \rangle = 0$ and $\langle \boldsymbol{\Theta}_i \boldsymbol{\Theta}_j \rangle = Q$.
The connection to the old problem is
$\mathcal{W} = Q^{-\top/2} W, \quad \boldsymbol{\Theta} = Q^{-\top/2} \Theta, \quad \mathcal{V} = Q^{-\top/2} V.$

Extremize

$$J := \langle \boldsymbol{\Theta}^\top \boldsymbol{\Theta} \rangle = \langle [Wx - V]^\top Q^{-1} [Wx - V] \rangle.$$

The Cholesky decomposition of $Q = Q^{\top/2} Q^{1/2}$.
Solve the Normal Equations $\mathcal{W}^\top \mathcal{W} \tilde{x} = \mathcal{W}^\top \mathcal{V}$, which yield

$$\begin{aligned} \tilde{x} &= (W^\top Q^{-1} W)^{-1} W^\top Q^{-1} V, && \text{the estimate,} \\ \tilde{n} = Q^{\top/2} \tilde{\mathbf{n}} &= (I - W(W^\top Q^{-1} W)^{-1} W^\top Q^{-1}) V, && \text{the residual,} \\ P &:= (W^\top Q^{-1} W)^{-1} W^\top Q^{-1} W (W^\top Q^{-1} W)^{-1}, && \text{uncertainty,} \end{aligned}$$

A common situation: $Q_{ij} = \mathbb{E}(\Theta_i \Theta_j)$.

Sequential Least Squares

Let $x(t) := [x_1, x_2]^\top$. Suppose you already have an estimate of x_1. Can we use this to find the estimate of x_2?

$$W_1 x_1 - V_1 = \Theta_1, \qquad W_2 x_2 - V_2 = \Theta_2.$$

Let $\langle \Theta_i \rangle = 0$, $\langle \Theta_i {\Theta_i}^\top \rangle = Q_i$. Assume additionally that $\langle \Theta_1 {\Theta_2}^\top \rangle = 0$. The global estimate is obtained by extremizing

$$J = \sum_{i=1}^{2} [W_i x_i - V_i]^\top Q_i^{-1} [W_i x_i - V_i].$$

Suppose we already have x_1 and P_1, then

$$\tilde{x}_2 = (W_1^\top Q_1^{-1} W_1 + W_2^\top Q_2^{-1} W_2)^{-1} (W_1^\top Q_1^{-1} V_1 + W_2^\top Q_2^{-1} V_2).$$

An expression can be written for $P_2 = \langle (x_2 - \tilde{x}_2)(x_2 - \tilde{x}_2)^\top \rangle$ as well.

However, using the *matrix inversion lemma*:
One can obtain

$$
\begin{aligned}
\tilde{x}_2 &= \tilde{x}_1 + K_2[V_2 - W_2\tilde{x}_1] \\
\tilde{n}_2 &= V_2 - W_2\tilde{x}_2 \\
P_2 &= P_1 - K_2 W_2 P_1 .
\end{aligned}
$$

$K_2 := P_1 W_2^\top [W_2 P_1 W_2^\top + Q_2]^{-1}.$

matrix inversion lemma,

$$
\begin{pmatrix} A & B \\ B^\top & C \end{pmatrix}
$$

where $A^\top = A$, $C^\top = C$, B rectangular and dimensionally commensurate. Then

$$
\begin{aligned}
(C - B^\top A^{-1} B)^{-1} &= C^{-1} - C^{-1} B^\top (B C^{-1} B^\top - A) B C^{-1} \\
A B^\top (C + B A B^\top)^{-1} &= (A^{-1} + B^\top C^{-1} B)^{-1} B^\top C^{-1}.
\end{aligned}
$$

Kalman Filter

Forecast

$$\begin{aligned} X^* &= MX(t)+B(t) \qquad t=0,1,\ldots, \\ P^* &= MP(t)M^\top \end{aligned}$$

in the absence of any other information, X^* is a state prediction, P^* is state uncertainty prediction. *Analysis*

$$\begin{aligned} X(t+1) &= X^*+K(t+1)[Y(t+1)-H(t+1)X^*], \\ P(t+1) &= P^*-K(t+1)H(t+1)P^* \end{aligned}$$

where the *Kalman Gain Matrix* is

$$K(t+1) := P^*H(t+1)^\top[H(t+1)P^*H(t+1)^\top+R(t+1)]^{-1}$$

$X(0)$ and $P(0)$ are known.

cf. Review in Jazwinski, Dover Pub

Estimation Using Perfect Models

Find the model parameters m such that

$$A(m)x - b = 0$$

m is the vector of parameters. Use field data

$$y = Hx + \varepsilon.$$

Cast as constrained optimization problem:

$$\min_{m,x} \frac{1}{2}||Hx - y||_C^2 + \beta \mathscr{R}(m)$$
$$\text{subject to } A(m)x - b = 0.$$

Estimation Using Perfect Models

Conventional Approach: *INCORPORATE CONSTRAINT*:

$$\min_m \frac{1}{2}||HA(m)^{-1}b - y||_C^2 + \beta\mathscr{R}(m).$$

Very compute-intensive:

- Each evaluation of the objective function requires a solution to the forward problem.
- Evaluating the gradient requires the solution to the adjoint problem.

Estimation Using Perfect Models

Alternative Approach: *ALL-IN-ONE OR AUGMENTED*:

$$\mathscr{L}(x,m,\lambda) = \frac{1}{2}||Hx - y||_C^2 + \beta \mathscr{R}(m) + \lambda^T (A(m)x - b).$$

The Euler-Lagrange equations are:

$$\begin{aligned}
\mathscr{L}_\lambda = A(m)x - b &= 0, \\
\mathscr{L}_x = A(m)^\dagger \lambda + H^\dagger (Hx - y) &= 0, \\
\mathscr{L}_m = \beta \frac{\partial \mathscr{R}}{\partial m} + \frac{\partial [A(m)x]}{\partial m}^\dagger \lambda &= 0.
\end{aligned}$$

Solve using Newton (preconditioned Krylov method). Same strengths-weaknesses of unconstrained method, but faster (only need approximate Hessian).

Estimation Using Non-Perfect Models

Find the model parameters m such that

$$A(m)x - b = \mu$$

m is the vector of parameters. Use field data

$$Hx - y = \varepsilon.$$

Known: $\mu \sim \mathcal{N}(0, C_\varepsilon)$ *and* $\varepsilon \sim \mathcal{N}(0, C_\mu)$.

Construct the over(under) determined system

$$W(m)x - V = \Theta.$$

Solve the weighted-row least-squares problem.

Model and Observations:

$$W(m)x - V = \Theta,$$
$$\Theta \sim \mathcal{N}(0, \sigma).$$

Find $\tilde{x}$ *, mean, such that* $\mathbb{E}(\theta^\top \theta)$ *is minimized.*
Find the uncertainty $U := \mathbb{E}[(x - \tilde{x})(x - \tilde{x})^\top]$.

Time Dependent Problems?

Consider a discrete time process...
Still can use Least Squares: suppose know x_0 and your model is

$$x_{n+1} = Mx_n + B_n + U_n + N_n,$$

$n = 0, 1, 2....$ Is your (discrete) linear time dependent model. Then it is easy to show that

$$x_n = Lx_0 + f(B_n) + g(U_n) + N$$

so we are back to solving a linear equation and can use Least Squares... but it might be more convenient to solve the estimation problem sequentially...

Time Dependent Problems

Task: *want to find* $\tilde{X}_n$*,* $n = 0, 1, ...$ *and uncertainty* P_n *that minimizes the posterior covariance of* X *at each* n*, given observations*

$$Y_n = HX_n + \varepsilon_n,$$

$n = 0, 1, 2....$ Here $\langle \varepsilon_n \rangle = 0$, $\langle \varepsilon_n \varepsilon_n^\top \rangle = R_n$. H is known as the observation matrix.

The model for the process is

$$X_{n+1} = MX_n + B_n + \Gamma U_n,$$

$n = 0, 1, 2....$ We assume $\langle U_n \rangle = 0$, $\langle U_n U_n^\top \rangle = Q_n$,

Note ΓU can be thought of as model noise (or it could be thought of as a CONTROLLER)

Kalman Filter, from $n = 0$ to $n = 1$:

Have an estimate of X_0, called $\tilde{X}_0$ with uncertainty P_0. The initial error is $\gamma_0 = \tilde{X}_0 - X_0$. *Forecast:* $X(1,-) = MX_0 + B_0$, The control (or noise) has zero mean and thus a best estimate is to set to zero. The -1 indicates that no data has been used in the estimate.

$$\gamma(1) = X(1,-) - X_1 = M\tilde{X}_0 + B_0 - (MX_0 + B_0 + \Gamma U_0) = M\gamma_0 - \Gamma U_0.$$

the erroneous forecast has 2 components: the propagated erroneous portion of $\tilde{X}_0$ and the unknown control term.

$$\langle \gamma_1 \gamma_1^\top \rangle = \langle (M\gamma_0 - \Gamma U_0)(M\gamma_0 - \Gamma U_0)^\top \rangle = MP_0M^\top + \Gamma Q_0 \Gamma^\top := P(1,-).$$

Use the measurement: $Y_1 = H_1 X_1 + N_1$: *Analysis*

$$\begin{aligned} \tilde{X}_1 &= X(1,-) + K_1[Y_1 - H_1 X(1,-)], \\ P_1 &= P(1,-) - K_1 H_1 P(1,-) \end{aligned}$$

where the *Kalman Gain Matrix* is $K_1 := P(1,-)H_1^\top [H_1 P(1,-) H_1^\top + R_1]^{-1}$

Kalman Filter Equivalent in Least Squares

The (sequential) Kalman filter estimate is also given by minimizing

$$
\begin{aligned}
E &= [X(1,-) - \tilde{X}_1]^\top P(1,-)^{-1} [X(1,-) - \tilde{X}_1] \\
&+ [Y_1 - H_1 \tilde{X}_1]^\top R_1^{-1} [Y_1 - H_1 \tilde{X}_1].
\end{aligned}
$$

Kalman Filter

Forecast

$$\begin{aligned} X^* &= MX_n + B_n \qquad n = 0, 1, \dots, \\ P^* &= MP_nM^\top + \Gamma Q_n \Gamma^\top \end{aligned}$$

Analysis

$$\begin{aligned} X_{(n+1)} &= X^* + K_{(n+1)}[Y_{(n+1)} - H_{(n+1)}X^*], \\ P_{(n+1)} &= P^* - K_{(n+1)}H_{(n+1)}P^* \end{aligned}$$

where the *Kalman Gain Matrix* is

$$K_{(n+1)} := P^* H^\top_{(n+1)}[H_{(n+1)}P^*H^\top_{(n+1)} + R_{(n+1)}]^{-1}$$

X_0 and P_0 are known.

Kalman Filter

The typical filter estimate, here observations have low variance:

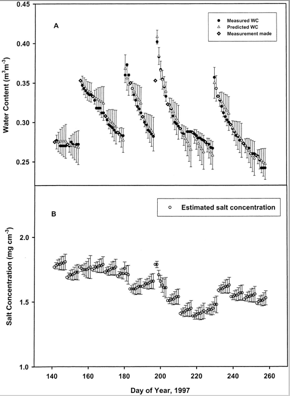

At filtering times there's a forecast correction due to the data (ANALYSIS). Between filtering times the uncertainty grows due to model errors.

Image from L. Wu, T. H. Skaggs, P. J. Shouse and J. E. Ayars *Soil Sci Soc. Amer* (2000)

Example: Feature Tracking

(Loading breakmovie)

- Uses 60 mpg frames of a basketball bouncing. (Data is (2d) edge of b'ball, found by edge detection)
- First order regression equation for the model.

Green is data and Red is the Extended Kalman Filter Estimate

taken from Mathworks, Inc, created Ali Reza Kashanipour

Example: Forced Coupled Oscillators

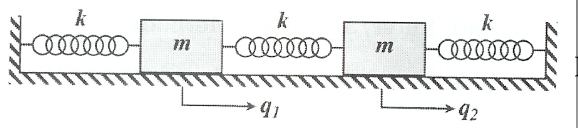

$$\mathbf{M}\frac{d^2\mathbf{q}}{dt^2} + \mathbf{R}\frac{d\mathbf{q}}{dt} + \mathbf{L}\mathbf{q} = \mathbf{f}$$

$$X_{n+1} = AX_n + F$$

where $X = [q_1\ q_2\ p_1\ p_2]^\top$, and

$$A = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ -2\alpha_1 & \alpha_2 & \beta_1 & 0 \\ \alpha_2 & -2\alpha_1 & 0 & \beta_2 \end{bmatrix}$$

$\alpha_i = dt\,k/m_i$, and $\beta_i = 1 - dt\,r_i/m_i$, $i = 1,2$. Also $F = [0\,0\,f_1/m_1\,f_2/m_2]^\top$.

Kalman Filter Problem

Estimate Mean Position and Uncertainty of Masses

- The forcing terms are ”noisy” and give the experiment some uncertainty in the observations.
- Observations of the position were made with a noisy device
- The goal is to use the model and the partial observations of the position of the masses to produce a filtered estimate of the vector $X = [p_1\ p_2\ q_1\ q_2]^\top$
- We will vary the measurement uncertainty, the frequency at which we sample the position

Conditions for the Experiment

- time step: 0.01
- number of time steps: 2000
- variance on forcing of q_1=0.75
- observed q_2 at every 15th time step

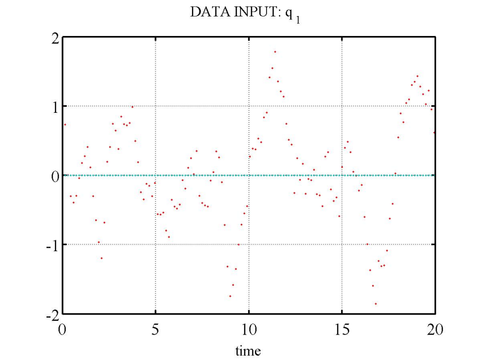

Results

- variance in measurement error = 0.05
- variance in model error = 0.08
- Compare norm of difference between *truth* to *filtered* as well as *truth* to *unfiltered*:
  - Maximum L_2 in *filtered* position of q_1 and $q_2 = 27$
  - Maximum L_2 in *unfiltered* position of q_1 and $q_2 = 37$

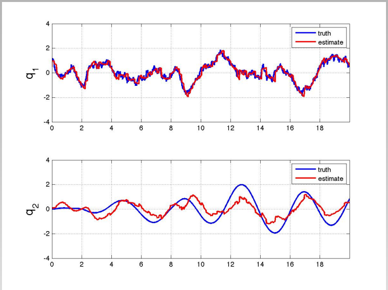

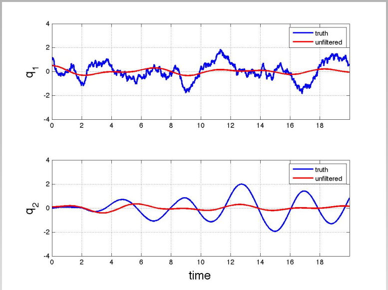

Results

- variance in measurement error = 0.05
- variance in model error = 0.08
- Compare norm of difference between *truth* to *filtered* as well as *truth* to *unfiltered*:
  - Maximum L_2 in *filtered* position of q_1 and $q_2 = 24.7$
  - Maximum L_2 in *unfiltered* position of q_1 and $q_2 = 42.3$

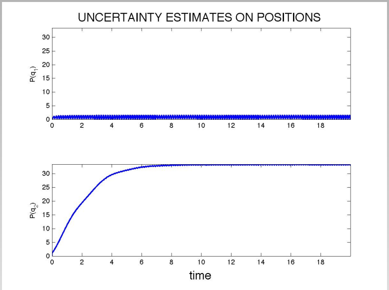

4D VAR: The Variational Approach

Goal: Find a posterior variance minimizer estimate $\hat{u}(x,t)$ *of the mean trajectory of* $u(x,t)$*, which obeys a noisy PDE and a noisy discrete data set* d_m

$$P(u(x,t)|d_{m=1:M}) \propto Likelihood \times Prior.$$

- Model informs prior,
- data informs the likelihood
- Assume (data and model) erros are normal, delta-correlated, with known variance.

The (Strong) Problem

MODEL:

$$
\begin{aligned}
Gu(x,t) &= F(x,t) + f(x,t), \quad 0 \le x \le L, 0 \le t \le T, \\
u(x,0) &= I(x) + i(x), \quad 0 \le x \le L \\
u(0,r) &= B(t) + b(t), \quad 0 \le t \le T,
\end{aligned}
$$

DATA:

$$
d_m = u(x_m, t_m) + \varepsilon_m, \quad m = 1, 2, .., M.
$$

where $G := \partial_t + c\partial_x$, and $c > 0$
$f(x,t), i(x), b(t), \varepsilon_m$ are normal noise processes with known variances:

$$
\langle f(x,t) f(x',t') \rangle = W_f^{-1}, \langle i(x) i(x') \rangle = W_i^{-1}, \langle b(t) b(t') \rangle = W_b^{-1}, \quad \langle \varepsilon_m \varepsilon_m' \rangle = w^{-1}.
$$

The Variational Problem

Let

$$\begin{aligned} J(u) &= W_f \int_0^T dt \int_0^L dx\, f(x,t)^2 + W_i \int_0^L dx\, i(x)^2 + W_b \int_0^T dt\, b(t)^2 \\ &+ w \sum_{m=1}^{M} \varepsilon_m^2 \delta \end{aligned}$$

where $\delta := \delta(x - x_m)\delta(t - t_m)$

$$J(\hat{u} + \delta u) = J(\hat{u}) + \mathscr{O}(\delta u^2),$$

since we force $\delta J(\hat{u}) = 0$.

$$
\begin{aligned}
0 = \delta J(\hat{u}) &= W_i \int_0^L dx\, [\hat{u}(x,0) - I(x)]\delta u(x,0) \\
&+ W_b \int_0^T dt\, [\hat{u}(0,t) - B(t)]\delta u(0,t) \\
&+ \int_0^L dx \int_0^T dt\, [-G\lambda]\delta u(x,t) \\
&+ \int_0^L dx\, \lambda\, \delta u|_{t=0}^{T} + \int_0^T dt\, c\lambda\, \delta u(x,t)|_{x=0}^{L} \\
&+ w \int_0^L dx \int_0^T dt \sum_{m=1}^{M} [\hat{u}(x,t) - d_m]\delta u(x,t)\delta
\end{aligned}
$$

where $\lambda = W_f(G\hat{u} - F)$.

The Euler-Lagrange Equations

With $\lambda = W_f(G\hat{u} - F)$,
BACKWARD PROBLEM

$$
\begin{aligned}
-G\lambda &= -\sum_{m=1}^{M} [\hat{u}(x,t) - d_m]\delta \\
\lambda(x,T) &= 0, \qquad \lambda(L,t) = 0,
\end{aligned}
$$

FORWARD PROBLEM

$$
\begin{aligned}
G\hat{u} &= F + W_f^{-1}\lambda \\
\hat{u}(x,0) &= I(x) + W_i^{-1}\lambda(x,0), \qquad \hat{u}(0,t) = B(t) + cW_b^{-1}\lambda(0,t).
\end{aligned}
$$

The best estimates of f, i, b:

$$
\hat{f}(x,t) = W_f^{-1}\lambda(x,t), \quad \hat{i}(x) = W_i^{-1}\lambda(x,0), \quad \hat{b}(t) = cW_b^{-1}\lambda(0,t).
$$

The Representer and the Reproducing Kernel

Let $r_m(x,t)$ and $\alpha_m(x,t)$ be the $m = 1 : M$ representers and adjoints,
ADJOINT PROBLEM:

$$\begin{aligned} -G\alpha_m &= \delta(x - x_m)\delta(t - t_m), \\ \alpha_m(x,T) &= 0, \quad \alpha_m(L,t) = 0 \end{aligned}$$

FORWARD PROBLEM:

$$\begin{aligned} Gr_m &= W_f^{-1}\alpha_m, \\ r_m(x,0) &= W_i^{-1}\alpha(x,0), \quad r_m(0,x) = cW_b^{-1}\alpha_m(0,t). \end{aligned}$$

ADJOINT PROBLEM:

$$\begin{aligned} -G\alpha_m &= \delta(x - x_m)\delta(t - t_m), \\ \alpha_m(x, T) &= 0, \quad \alpha_m(L, t) = 0 \end{aligned}$$

FORWARD PROBLEM:

$$\begin{aligned} Gr_m &= W_f^{-1}\alpha_m, \\ r_m(x, 0) &= W_i^{-1}\alpha(x, 0), \quad r_m(0, x) = cW_b^{-1}\alpha_m(0, t). \end{aligned}$$

$$\hat{u} = u_F(x, t) + \sum_{m=1}^{M} \beta_m r_m(x, t)$$

Need to find β_m's in

$$\hat{u} = u_F(x,t) + \sum_{m=1}^{M} \beta_m r_m(x,t).$$

Substitute $\hat{u}$ into the *forward problem* equation $G\hat{u} = F + W_f^{-1}\lambda(x,0)$, to find

$$G\hat{u} = Gu_F + \sum_{m=1}^{M} \beta_m G r_m = F + W_f^{-1} \sum_{m=1}^{M} \beta_m \alpha_m.$$

Thus

$$\lambda = W_f[G\hat{u} - F] = \sum_{m=1}^{M} \beta_m \alpha_m.$$

Further, using $-G\lambda = -w\sum_{m=1}^{M}[\hat{u}(x,t)-d_m]\delta(x-x_m)\delta(t-t_m)$ from the *backward problem*,

$$
\begin{aligned}
-G\lambda &= -\sum_{m=1}^{M}\beta_m G\alpha_m = \sum_{m=1}^{M}\beta_m\delta(x-x_m)\delta(t-t_m) = \\
&- \quad w[\hat{u}(x,t)-d_m]\delta(x-x_m)\delta(t-t_m).
\end{aligned}
$$

Which implies

$$
\beta_m = -w[\hat{u}(x,t)-d_m]\delta(x-x_m)\delta(t-t_m).
$$

Substituting $\hat{u} = u_F(x,t) + \sum_{m=1}^{M}\beta_m r_m(x,t)$,

$$
\beta_m = -w[u_F(x_m,t_m) + \sum_{\ell=1}^{M}\beta_\ell r_\ell(x_m,t_m) - d_m].
$$

Hence,

$$
\sum_{\ell=1}^{M}[r_\ell(x_m,t_m) + w^{-1}\delta_{\ell,m}]\beta_\ell = d_m - u_F(x_m,t_m).
$$

The best estimate is

$$\hat{u} = u_F(x,t) + \sum_{m=1}^{M} \beta_m r_m(x,t),$$

where

$$\sum_{\ell=1}^{M} [r_\ell(x_m,t_m) + w^{-1}\delta_{\ell,m}]\beta_\ell = d_m - u_F(x_m,t_m),$$

or

$$[\mathbb{R} + w^{-1}\mathbb{I}]\beta = \mathbf{d} - \mathbf{u_F},$$

Finally:

$$\hat{u}(x,t) = u_F(x,t) + (\mathbf{d} - \mathbf{u_F})^\top [\mathbb{R} + w^{-1}\mathbb{I}]^{-1}\mathbf{r}(x,t).$$

Nonlinear Non-Gaussian Problems?

Forecast, not much of a problem:

$$\tilde{X} = N(X(t), B(t))$$

But not clear how to propagate uncertainty $P(t+1)$.

Extended Kalman Filter used extensively on nonlinear problems: linearize about $X(t)$ and use closure ideas for moments.

The EKF Results[1]

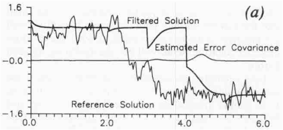

Figure: 10% uncertainty, $\Delta t = 1$.

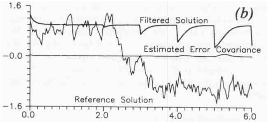

Figure: 20% uncertainty, $\Delta t = 1$.

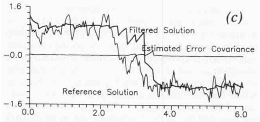

Figure: 20% uncertainty, $\Delta t = 0.25$.

[1] R. Miller, M. Ghil, P. Gauthiez, *Advanced data assimilation in strongly nonlinear dynamical systems*, J. Atmo. Sci. **51** 1037-1056 (1994)

Rao Blackwellisation: Reduce Variance

An essential dimensional reduction stage: identify linear/Gaussian elements in your state vector and use Kalman (or least squares) on these....it's optimal!

Rewrite $x = x^l, x^n$, then

$$p(X|Y) \propto p(x^l|x^n, Y)p(x^n|Y).$$

use your nonlinear/non-Gaussian sampler on $p(x^n|Y)$.

$$var(\mathbb{E}[g(x^l|x^n)|x^n)] + \mathbb{E}[var(g(x^l, x^n)|x^n)] = var(g(x^l|x^n))$$

thus, $var(\mathbb{E}g(x^l|x^n)]|x^n)] \leq var(g(x^l|x^n))$.

cf., See Karlsson, Sh on, Gustaffson, IEEE Trans. Sig, Proc. 2005

Other Approaches on Nonlinear/Non-Gaussian Problems

- Optimal (variance-minimizer) KSP (Kushner, Stratonovich, Pardoux), early 60's
- 4D-Var/Adjoint (Maximum Likelihood) (Wunsch, McLaughlin, Courtier, late 80's)
- ensemble KF (Evensen, '97)
- Mean Field Variational (Rayleigh-Ritz on the Kullback-Leibler Divergence) (Eyink, Restrepo, '01)
- Parametrized Resampling Particle Filter (Kim, Eyink, Restrepo, Alexander, Johnson, '02)
- Langevin Sampler (A. Stuart, '05)
- Path Integral Monte Carlo (Restrepo '07. Alexander, Eyink & Restrepo, '05)
- Diffusion Kernel Filter (Krause, Restrepo, '09)
- Displacement Assimilation (Venkataramani, Rosenthal, Mariano, Restrepo, '13)
- Mean Stochastic Sampler (Harlim and Majda, '10)

Restrepo, Leaf, Griewank, *Circumventing storage limitations in variational data assimilation*, SIAM J. Sci Comp, '95

enKF Most Favored in Practice

The enKF (”state-of-the-art”)

- Use model for forecast $\tilde{X} = N(X(t), t)$.
- Update the uncertainty using Monte Carlo.

Pros and Cons:

- Can handle legacy code easily
- *Gaussian assumption on the analysis*: $X(t+1) = \tilde{X} + K(t)(y - H(\tilde{X}))$.
- Requires full model runs
- Ad-hoc

G. Evensen, Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods to forecast error statistics, J. Geophys. Res. **99**, 10143-10162.

PIMC The Path Integral Monte Carlo

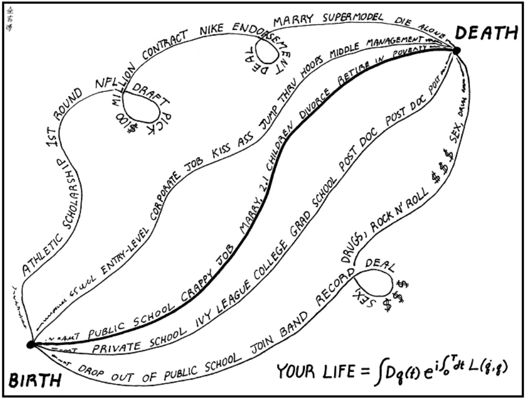

The Path Integral Formulation of Your Life

J. Restrepo, *A Path Integral Method for Data Assimilation*, Physica D, 2007,
F. Alexander, G. Eyink, J. Restrepo, *Accelerated Monte-Carlo for Optimal Estimation of Time Series*, J. Stat. Phys., 2005

PIMC The Path Integral Monte Carlo

- Optimal, on the discretized model
- Simple to implement, *but very subtle*
- Can handle legacy code
- Relies on sampling
- Can yield a variety of different estimators

J. Restrepo, A Path Integral Method for Data Assimilation, Physica D, 2007,
F. Alexander, G. Eyink, J. Restrepo, Accelerated Monte-Carlo for Optimal Estimation of Time Series, J. Stat. Phys., 2005

PIMC The Path Integral Monte Carlo

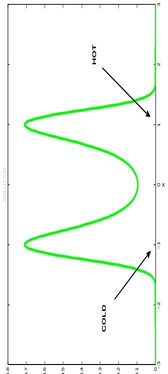

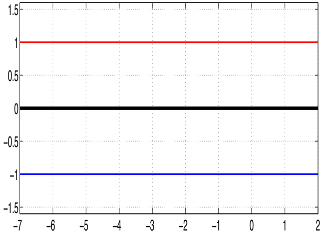

PIMC The Path Integral Monte Carlo

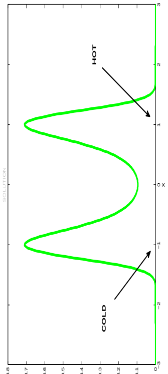

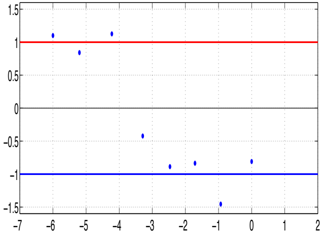

PIMC The Path Integral Monte Carlo

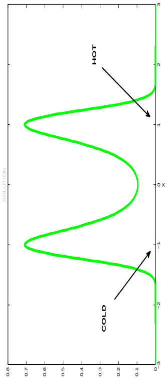

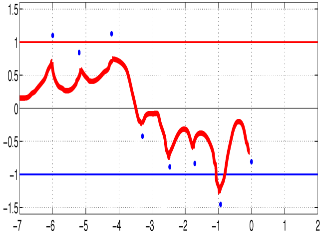

PIMC The Path Integral Monte Carlo

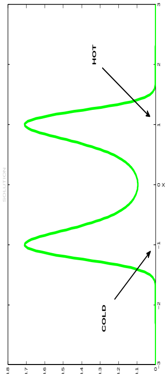

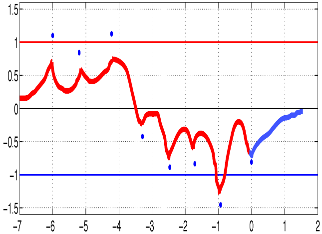

Bayesian Statement

- $P(x|y) \propto$ Likelihood $\times$ Prior.
- Use data for likelihood.
- Use model for prior.

$$P(x|y) \propto e^{-\mathscr{A}_{model}} e^{-\mathscr{A}_{data}} := e^{-\mathscr{A}(x)}.$$

$\mathscr{A}_{model}$

$$dx = f(x,t)dt + [2D(x,t)]^{1/2} dW$$

is discretized:

$$x_{n+1} = x_n + \Delta t f(x_n, t_n) + [2D(x_n, t_n)]^{1/2} [W_{n+1} - W_n]$$

$$n = 0, 1, ..., T-1$$

$$\mathscr{A}_{model} \approx \sum_{n=1}^{T} [(x_{n+1} - x_n - \Delta t f(x_n, t_n))^\top D(x_n, t_n)^{-1} (x_{n+1} - x_n - \Delta t f(x_n, t_n))],$$

if $\text{Prob}(\Delta W) \propto exp(-\Delta W^2/D)$.

$\mathscr{A}_{data}$

$$y_m = H(x_m) + [2R[x_m, t_m)]^{1/2}\eta_m$$
$$m = 1, 2, \ldots, M.$$

$$\mathscr{A}_{data} = \sum_{m=1}^{M} [(y_m - H(x_m))^\top R(x_n, t_n)^{-1} (y_m - H(x_m))],$$

if $\text{Prob}(\eta) \propto exp(-\eta^2/R)$.

MCMC Samplers

$$P(x|y) \propto e^{-\mathscr{A}_{model}} e^{-\mathscr{A}_{data}} := e^{-\mathscr{A}(x)}.$$

The Path Integral Monte Carlo practicality depends on fast sampling:

- Multigrid (UMC)
- Langevin Sampler (LS)
- Hybrid Monte Carlo (HMC)
- Shadow Hybrid MC (sHMC)
- Riemannian Manifold Hamiltonian Monte Carlo (RM-HMC)
- generalized Hybrid Monte Carlo (gHMC)

(HMC) Hybrid Markov Chain Monte Carlo

- Proposals generated by solving Hamiltonian system in fictitious time τ.
- Accept/reject via Metropolis Hastings

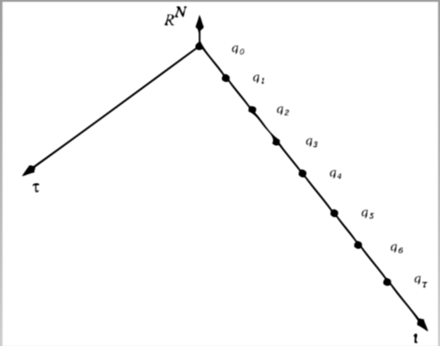

HMC Algorithm

Let $q_n(\tau = 0) = x_n$.

- To each q_n, a conjugate generalized momemtum, p_n, is assigned.
- The momenta p_n give rise to a kinetic contribution $K = \sum_{n=1}^{T} p_n^\top M^{-1} p_n / 2$.
- The Hamiltonian of the system $\mathscr{H} = \mathscr{A}(q) + K(p)$.
  The dynamics are:

$$\begin{aligned} \frac{\partial q_n}{\partial \tau} &= M^{-1} p_n \\ \frac{\partial p_n}{\partial \tau} &= F_n \quad \text{where} \quad F_n = -\text{grad}(\mathscr{A}(q)). \end{aligned}$$

- Solve using Verlet integrator (detailed balance).
- Accept/Reject Metropolis/Hastings.

Why does HMC work? What are good HMC properties?

Write probability $\Pi(q) = \frac{1}{Z_\Pi} e^{-\mathscr{A}(q)}$:

- Sampling $\pi(q,p) = \frac{1}{Z_\pi} e^{-\mathscr{H}(q,p)} \sim \frac{1}{Z} e^{-\mathscr{A}(q)}$ samples $\Pi(q)$.
- Gradient dynamics makes system search through configuration space more efficiently.
- Moves in q_n are linear in p_n, *i.e.*, $\frac{\partial q}{\partial \tau} = M^{-1} p$

$\mathscr{A}(q)$ and grad$(\mathscr{A}(q))$ should be easily evaluated.

Sampler Efficiency Estimates

Sampler Efficiency: key to choosing and tuning sampler

- Computational Cost: $\mathscr{O}(NT)^r n_{\text{method}}(p,L)$
- $p := < P_{acc} > = < \min\{1, \exp[-\Delta\mathscr{H}]\} > \propto \text{erfc}\left(\frac{1}{2}\delta\tau^m (NT)^{1/2}\right)$.
- $c(L) := < \mathscr{H}(0)\mathscr{H}(0+L) >$. Depends on problem dimension and state space characteristics.

RM-HMC Algorithm[2]

Hamiltonian replaced by:

$$\mathscr{H} = \mathscr{A}(q) + \frac{1}{2} p^\top G(q)^{-1} p$$

where the *non-degenerate Fisher information matrix* $G := \mathbb{E}\{\nabla\mathscr{A}\nabla\mathscr{A}^\top\}$

Challenges:

- find a time-reversible/volume-preserving discrete integrator for Hamiltonian problem.
- optimize its computational efficiency.

[2] Girolami, Calderhead, Chin, preprint, 2009.

Decrease decorrelation length L: gHMC Algorithm

Hamiltonian dynamics replaced by:

$$\frac{\partial q_n}{\partial \tau} = CM^{-1}p_n$$
$$\frac{\partial p_n}{\partial \tau} = C^\top F(q_n)$$

where $C \in \mathbb{R}^{T\times T}$ matrix

Challenge: find C that leads to a significant reduction in the sample decorrelation length.

We used the circulant matrix $C = \text{circ}(1, e^{-\alpha}, e^{-2\alpha}, \ldots, e^{-T\alpha})$.

Sampler Efficiency Comparison

Table: *T is the number of time steps, $(\cdot)$ is the standard deviation on the number of samples, $[\alpha]$ used in C; J is the number of τ time steps.*

$T+1$	HMC (J=1)	HMC (J=8)	UMC	gHMC (J=1)
8	900(125)	170(7)	800(40)	40(8) [0.20]
16	5300(1600)	560(20)	1040(60)	60(10) [0.10]
32	13300(8300)	2700 (140)	1430(100)	200(30) [0.05]
64	30000(7800)	2800(400)	1570(100)	420(70) [0.0245]

Looking Forward...

- Continued work on assimilation methods that can handle larger problems.
- Continue improving nonlinear/non-Gaussian assimilation methods.

Data and models can combine to improve forecasts...but can they be used to make better forecasts?

- Feature-based data assimilation.
- Displacement data assimilation.
- Surrogate models built from data only.

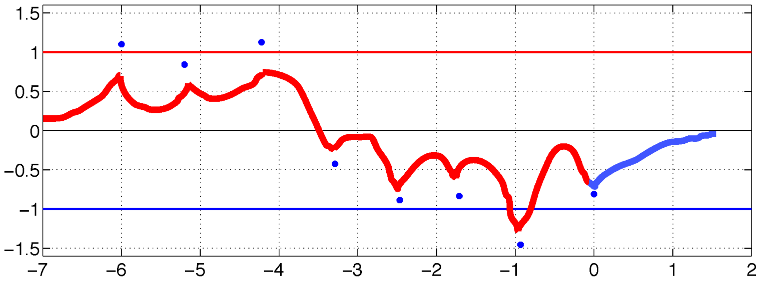
1.5
1
0.5
0
−0.5
−1
−1.5
−7
−6
−5
−4
−3
−2
−1
0
1
2

Feature-Based, Lagrangian Data Blending

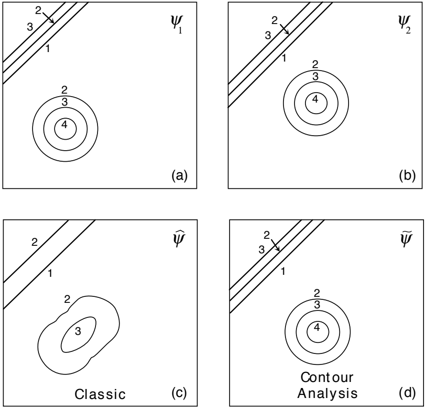

Improving Hurricane Predictions

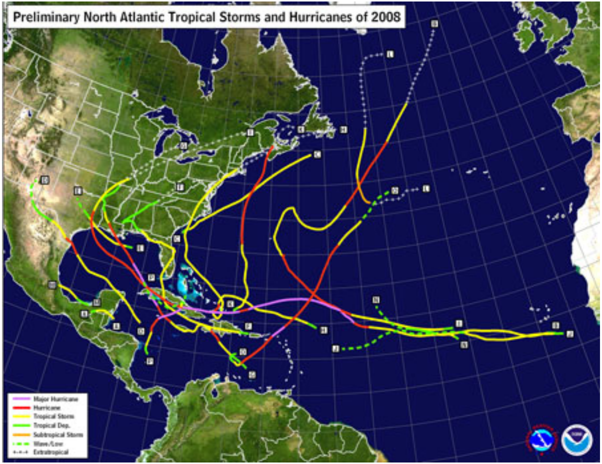

Improving Hurricane Predictions

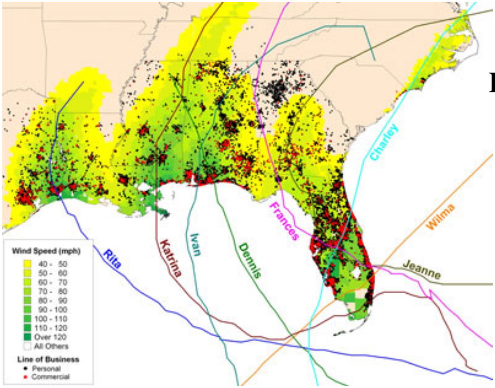

Estimated Property Damage

- Katrina $108B
- Sandy $65B
- Ike $30B
- Andrew $27B
- ...

Goal: Better Predictions using Added Constraints

- Better estimates of hurricane tracks (NSF)
- Better estimates of oil slick geometry and location in ocean flows (BP/GoMRI)

Collaborators:
Steven Rosenthal (Arizona)
Shankar Venkataramani (Arizona)
Arthur Mariano (U Miami)

Displacement Maps via Canonical Transformations

Find

$$\min ||q(M(x)) - q_0||_2^2.$$

here $(x,y) \xrightarrow{M} (X,Y)$.
In 2-Dimensions, the generating function is $G(X,y) = Xy + f(X,y)$.

$$x = \frac{\partial G}{\partial y} = X + f_y(X,y)$$
$$Y = \frac{\partial G}{\partial X} = y + f_X(X,y).$$

invertible if $f_{yX} > -1$.

Parameterizing Position Error

Example:

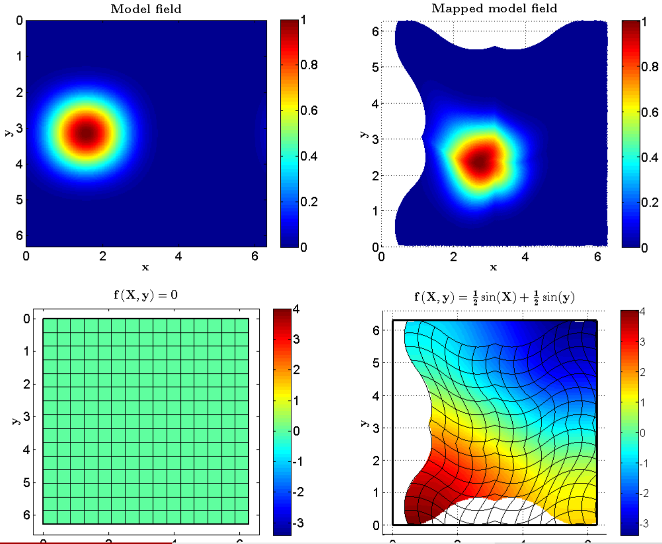

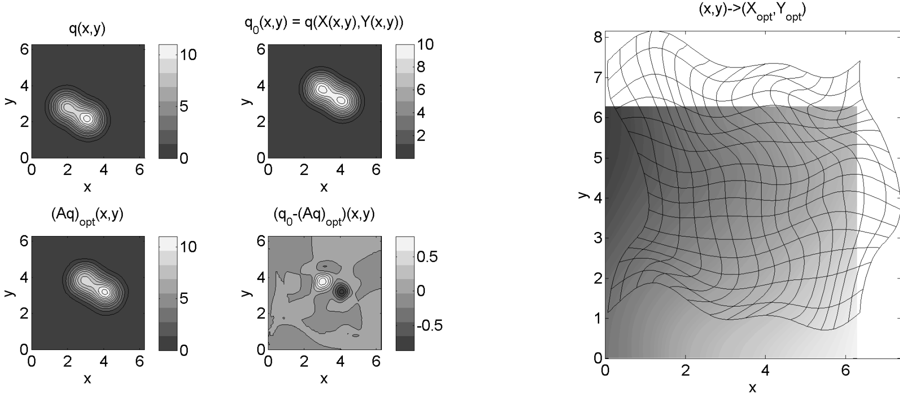
q(x,y)
q0(x,y) = q(X(x,y),Y(x,y))
(Aq)opt(x,y)
(q0-(Aq)opt)(x,y)
(x,y)->(Xopt,Yopt)
x
y

The strain tensor σ takes the form

$$\sigma = \begin{bmatrix} x\Delta x & y\Delta x \\ x\Delta y & y\Delta y \end{bmatrix} = \frac{1}{1+f_{yX}} \begin{bmatrix} -f_{yX} & -f_{yy} \\ f_{XX} & f_{Xy} - |H[f]| \end{bmatrix}$$

where $H[f]$ is the Hessian matrix of f. The diagonal terms determine the normal strains in the map, while the off-diagonal terms define the shear strains. The penalty functional is now given by

$$\mathscr{J}[f] = \int_D [q(f) - q_0]^2 + \alpha \left[(x\Delta x)^2 + (y\Delta y)^2 \right] + \beta \left[(y\Delta x)^2 + (x\Delta y)^2 \right] \mathrm{d}x\,\mathrm{d}y$$

where α and β weight the normal and shear strains, respectively.

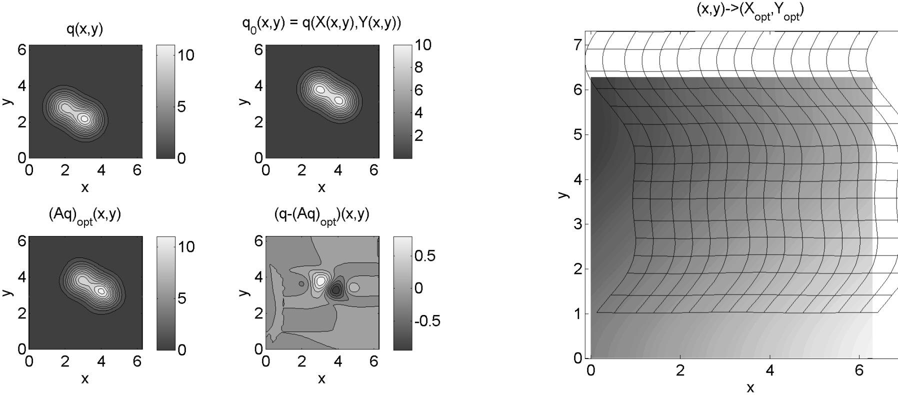
q(x,y)
q0(x,y) = q(X(x,y),Y(x,y))
(Aq)opt(x,y)
(q-(Aq)opt)(x,y)
(x,y)->(Xopt,Yopt)
x
y

Displacement and Amplitude Assimilation

Combine ”Traditional” Amplitude Assimilation with Displacement Assimilation:

Basic Algorithm (from t_m to t_{m+1}):

- At t_m: Perform displacement assimilation.
- At t_{m+1}: Perform amplitude assimilation.

Further Information

Juan M. Restrepo

www.physics.arizona.edu/~restrepo

Uncertainty Quantification Group

www.physics.arizona.edu/~restrepo/UQ.html