Partial Differential Equations (PDEs)

1. Classification

(a) Many PDEs encountered by chemical engineers are second order (containing at most second derivatives) or can be decomposed into a system of second order PDEs. The standard form for a linear second order PDE for a function $f : \mathbb{R}^N \to \mathbb{R}$ and independent variables $\mathbf{y} \in \mathbb{R}^N$ is

$$\sum_{i=1}^{N}\sum_{j=1}^{N} A_{ij}\frac{\partial^2}{\partial y_i \partial y_j}f(\mathbf{y}) + \sum_{i=1}^{N} b_i(\mathbf{y})\frac{\partial}{\partial y_i}f(\mathbf{y}) + g(\mathbf{y})f(\mathbf{y}) = h(\mathbf{y})$$

where $\mathbf{A}(\mathbf{y}) \in \mathbb{R}^{N \times N}$, $\mathbf{b} \in \mathbb{R}^N$, $g(\mathbf{y}) \in \mathbb{R}$, and $h(\mathbf{y}) \in \mathbb{R}$ are arbitrary functions of the independent variable $\mathbf{y}$. This equation is valid for some points $\mathbf{y}$ within some prescribed domain $\Omega \subset \mathbb{R}^N$, which has a boundary $\delta\Omega \subset \mathbb{R}^{N-1}$.

(b) Second order PDEs are classified by examining the properties of $\mathbf{A}(\mathbf{y})$ and $\mathbf{b}(\mathbf{y})$.

   i. **Elliptic** at point $\mathbf{y}$ if $\mathbf{A}(\mathbf{y})$ is not singular and all of its eigenvalues have the same sign.

   ii. **Hyperbolic** at point $\mathbf{y}$ if $\mathbf{A}(\mathbf{y})$ is not singular and all of its eigenvalues but one have the same sign.

   iii. **Parabolic** at point $\mathbf{y}$ if $\mathbf{A}(\mathbf{y})$ has one eigenvalue equal to zero and all other eigenvalues have the same sign. Additionally, the rank of the matrix

$$\begin{bmatrix} \mathbf{A}_1^C(\mathbf{y}) & \mathbf{A}_2^C(\mathbf{y}) & \cdots & \mathbf{A}_N^C(\mathbf{y}) & \mathbf{b}(\mathbf{y}) \end{bmatrix}$$

   must be N. In other words, $\mathbf{b}$ cannot be written as a linear combination of the columns of $\mathbf{A}(\mathbf{y})$.

The PDE is called an elliptic, hyperbolic, or parabolic PDE when these conditions hold for all $\mathbf{y}$ in the domain Ω.

(c) Quintessential examples of these types

   i. **Elliptic** – Steady diffusion represented by Laplace's equation

$$0 = D\left(\frac{\partial^2 C}{\partial x_1^2} + \frac{\partial^2 C}{\partial x_2^2}\right) \quad \xrightarrow[\mathbf{y}=(x_1,x_2)]{} \quad \mathbf{A}(\mathbf{y}) = \begin{bmatrix} D & 0 \\ 0 & D \end{bmatrix}$$

   ii. **Hyperbolic** – Wave equation

$$\frac{\partial^2 C}{\partial t^2} = u^2\left(\frac{\partial^2 C}{\partial x_1^2} + \frac{\partial^2 C}{\partial x_2^2}\right) \quad \xrightarrow[\mathbf{y}=(x_1,x_2,t)]{} \quad \mathbf{A}(\mathbf{y}) = \begin{bmatrix} u^2 & 0 & 0 \\ 0 & u^2 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

iii. **Parabolic** – Unsteady diffusion

$$\frac{\partial C}{\partial t} = D\left(\frac{\partial^2 C}{\partial x_1^2} + \frac{\partial^2 C}{\partial x_2^2}\right) \quad \xrightarrow[\mathbf{y}=(x_1,x_2,t)]{} \quad \mathbf{A}(\mathbf{y}) = \begin{bmatrix} D & 0 & 0 \\ 0 & D & 0 \\ 0 & 0 & 0 \end{bmatrix} ; \ \mathbf{b} = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}$$

2. Finite differences

   (a) In the finite difference method, the domain of a function $f(\mathbf{y})$ described by a PDE is decomposed into a finite set of points (or nodes). Then, the finite difference formula is used to approximate derivatives in the PDE at these nodes.

   (b) These difference formulas introduce errors with respect to spacing between the nodes. Typically, we use approximations that are first or second order with respect to the spacing between nodes.

   (c) With one discrete algebraic equation produced per node and $f(\mathbf{y})$ unknown at each node, this procedure yields a complete system of algebraic equations that can be solved using standard methods. Linear PDEs will result in a linear system of equations while nonlinear PDEs will result in a nonlinear system of equations.

   (d) This method can be used to solve most PDE problems (elliptic, parabolic, hyperbolic) including steady state or unsteady problems. It is fairly easy to implement but hard to incorporate irregular boundary shapes. This method does not conserve, conserved quantities (loss of mass, energy, momentum from simulation for example).

3. Finite volumes

   (a) The finite volume method was discussed in the context of BVPs. The idea does not change when we move to PDEs except now we look at cells possibly distributed in space and as a function of time.

   (b) The basic idea is to integrate the differential form of the conservation equation and then look at the "flux" (total amount of material) through each face of the cell. This allows us to look at very complicated geometries in a simple way as we are merely breaking up the domain into "cells" and satisfying a total conservation equation locally in each cell.

   (c) Typically the flux is approximated using some polynomial interpolation between cell centers (that represent the average concentration in the cell). See page 101 of the extended PDE notes for detailed equations.

   (d) The same rules apply for indexing the cells as discussed for nodes in the finite difference section. Again, the equations tend to be sparse such that we can take advantage of sparse/iterative solvers to get reasonably fast solutions.

4. Numerical method of lines

   (a) The method of lines is a technique for solving PDEs with initial conditions that capitalizes on the efforts to develop stable, high accuracy methods for first order ODEs.

   (b) Typically, the spatial domain is approximated using finite differencing while time derivatives are kept. This gives us a set of ODE-IVPs that can be solved using standard methods. This allows us to overcome issues with the advection equation using central difference in space for example (which is second order accurate) . We showed that a

forward Euler integration of this scheme was unconditionally unstable. However, when the time derivative is kept intact, ODE solvers that have a conditional stability as a result of adaptive time stepping or high order accuracy can be used.

Probability Theory

1. Preliminaries: probability obeys three rules:

   (a) For any event A, the probability of event A satisfies $\Pr(A) \geq 0$.

   (b) The probability of the event consisting of all possible outcomes is one, $\Pr(E) = 1$.

   (c) If the events A and B are mutually exclusive, then $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B)$.

   Any other result in probability theory, intuitive or otherwise, can be derived from these three rules without any further assumptions. One fundamental result that follows from these properties is: For any events A and B,

   $$\Pr(A \text{ or } B) = \Pr(A) + \Pr(B) - \Pr(A \text{ and } B)$$

2. Definitions

   (a) **Random variable** – Formally, a random variable is any function of the outcome of an experiment $g : E \to \mathbb{R}$. In general, let x denote a random variable on a set of outcomes E. Performing a single trial produces an outcome $\xi \in E$ that is a *realization* or *sample* of the random variable x, $x(\xi) = \hat{x} \in \mathbb{R}$. Since the value of a random variable depends on the outcome of a trial, the events can be defined in terms of random variables. For example, the subset of E defined by $\{\xi : x(\xi) \geq 5\}$ is an event, which is interpreted as *"the outcome of the trial ξ satisfies $\xi \geq 5$"*. These events also have probabilities, which are denoted by

   $$\Pr(x \geq 5) = \Pr(\{\xi : x(\xi) \geq 5\}).$$

   (b) **Probability mass function** – for a discrete random variable X, the probability mass of x is the probability of the event $X = x$ which can be written $p_X(x) = \mathbf{P}(X = x)$. A PMF must have the property that $\sum_x p_X(x) = 1$.

   (c) **Probability density function** – the probability density describes the probability that a continuous random variable takes a value in a specified range: $\mathbf{P}(a \leq X \leq b) = \int_a^b f_X(x)dx$. The PDF f_X satisfies

   $$\int_{-\infty}^{+\infty} f_X(x)dx = 1$$

   Because the PDF represents a "probability mass per unit length" the values of $f_X(x)$ do not need to be strictly less than 1 but they must be nonnegative.

   (d) **Cumulative probability function** – The cumulative probability function (also known as the cumulative distribution function) F_x of a random variable is defined as

   $$F_X(x) = \mathbf{P}(X \leq x) = \int_{-\infty}^x f_X(t)dt$$

F_x is a non-decreasing function with $F_x(-\infty) = 0$ and $F_x(+\infty) = 1$. The probability that x will be in the range $[x_1, x_2]$ in a single trial is

$$\Pr([x_1, x_2]) = F_x(x_2) - F_x(x_1)$$

(e) **Marginal probability** – given the joint probability $p_{X,Y}(x, y)$ the marginal probability of X is $p_X(x) = \sum_y p_{X,Y}(x, y)$

(f) **Conditional probabilities** – For two events A and B, the conditional probability of A given B is the probability that the event A will occur, given that event B is known to occur (or to have occurred). This is related to the intersection of the two sets by

$$\Pr(A|B) = \frac{\Pr(A \text{ and } B)}{\Pr(B)}$$

We can rearrange this expression for the conditional probability

$$\Pr(A \text{ and } B) = \Pr(A|B)\Pr(B)$$

(g) **Independence** The events A and B are said to be *independent* if

$$\Pr(A \text{ and } B) = \Pr(A)\Pr(B)$$

which implies that the events A and B are independent if and only if $\Pr(A|B) = \Pr(A)$. That is, the probability of the occurrence of event A is not affected by whether the event B has occurred. This result is consistent with our intuition as to the meaning of two events being independent.

(h) **Expected value** – The expected value of x (known as the *mean* of x) is

$$\langle x \rangle = \int_{-\infty}^{+\infty} x' f_X(x') dx'.$$

Let g be a function of the random variable x. The *expected value* of g is defined as

$$\langle g(x) \rangle = \int_{-\infty}^{+\infty} g(x') f_X(x') dx'.$$

(i) **Variance** – The *variance* of x, σ_x^2, is defined as the expected value of $(x - \langle x \rangle)^2$,

$$\sigma_x^2 = \langle (x - \langle x \rangle)^2 \rangle$$

$$= \int_{-\infty}^{+\infty} (x - \langle x \rangle)^2 p_x(x') dx'$$

$$= \int_{-\infty}^{+\infty} (x^2 - 2\langle x \rangle x + \langle x \rangle^2) p_x(x') dx'$$

$$= \int_{-\infty}^{+\infty} x^2 p_x(x') dx' - 2\langle x \rangle \underbrace{\int_{-\infty}^{+\infty} x p_x(x') dx'}_{\langle x \rangle} + \langle x \rangle^2 \underbrace{\int_{-\infty}^{+\infty} p_x(x') dx'}_{1}$$

$$= \langle x^2 \rangle - 2\langle x \rangle^2 + \langle x \rangle^2$$

$$= \langle x^2 \rangle - \langle x \rangle^2$$

The variance of $f(x)$, which is a function of random variable x, is

$$\sigma_f = \langle f^2 \rangle - \langle f \rangle^2$$

The *standard deviations* of x and f are $\sigma_x = \sqrt{\sigma_x^2}$ and $\sigma_f = \sqrt{\sigma_f^2}$, respectively.

3. Important Theorems

   (a) **Total probability theorem** – Let $A_1, \ldots, A_n$ be disjoint events that form a partition of the sample space (each possible outcome is included in exactly one of the events $A_1, \ldots, A_n$) and assume that $\mathbf{P}(A_i) > 0 \forall i$. Then, for any event B,

   $$\Pr(B) = \Pr(B|A_1)\Pr(A_1) + \Pr(B|A_2)\Pr(A_2) + \cdots + \Pr(B|A_n)\Pr(A_n)$$

   This can be written in a more compact form

   $$\Pr(B) = \sum_{i=1}^{n} \Pr(B \text{ and } A_i) = \sum_{i=1}^{n} \Pr(B|A_i)\Pr(A_i)$$

   (b) **Bayes' theorem** – For two events A_i and B, taking the intersection of two sets is independent of the order of the two sets,

   $$\Pr(A_i \text{ and } B) = \Pr(B \text{ and } A_i)$$

   Applying the conditional probability expression to both sides gives

   $$\Pr(A_i|B)\Pr(B) = \Pr(B|A_i)\Pr(A_i) \quad \longrightarrow \quad \Pr(A_i|B) = \frac{\Pr(B|A_i)\Pr(A_i)}{\Pr(B)}$$

   This is known as *Bayes' Theorem*. When A_i is a member of a collection of events, $A_1, \ldots, A_n$, which is mutually exclusive and collectively exhaustive, we can substitute the Total Probability Theorem into the denominator

   $$\Pr(A_i|B) = \frac{\Pr(B|A_i)\Pr(A_i)}{\Pr(B|A_1)\Pr(A_1) + \Pr(B|A_2)\Pr(A_2) + \cdots + \Pr(B|A_n)\Pr(A_n)}$$

   Bayes' theorem provides a useful framework for using experimental measurements to update knowledge about a physical problem, which is often applied in design of prognostic/diagnostic systems and in the estimation of states or model parameters from experimental data.

   (c) **Central limit theorem** – Let $X_1, X_2 \ldots$ be a squence of independent identically distributed random variables with common mean μ and variance σ^2, and define

   $$Z_n = \frac{X_1 + \cdots + X_n - n\mu}{\sigma\sqrt{n}}$$

   Then the CDF of Z_n converges to the standard normal CDF

   $$F(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{z} e^{-x^2/2} dx$$

   in the sense that $\mathbf{P}(Z_n \leq z) = F(z)$.
   $\underset{n\to\infty}{}$

4. Notable distributions

   (a) The *normal* or *Gaussian* PDF is

   $$p_x(\hat{x}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)}$$

   The mean and variance of a random variable x with this PDF are $\langle x \rangle = \mu$ and $\sigma_x^2 = \sigma^2$.

   (b) The *uniform* PDF on interval $[a, b] \subset \mathbb{R}$ is

   $$p_x(\hat{x}) = \begin{cases} \dfrac{1}{b-a}, & \text{if } \hat{x} \in [a, b] \\ 0, & \text{otherwise} \end{cases}$$

   A random variable x with this PDF has zero probability of being observed outside of the interval $[a, b]$, and the probability of being observed in any infinitesimal interval $[\hat{x}, \hat{x} + dx]$ is equal for all $\hat{x} \in [a, b)$. The mean and standard deviation of the uniform distribution are $\langle x \rangle = \frac{a+b}{2}$ and $\sigma_x = \frac{1}{3}\left(\frac{b-a}{2}\right)^2$.

   (c) The *binomial* PMF is

   $$p(k) = \binom{n}{k} p^k (1-p)^{n-k}$$

   where

   $$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

   and this factor is known as the binomial coefficient. When $n = 1$ this is known as the Bernoulli distribution.

5. Multivariate extensions

   (a) **Random vector** – A *random vector* is a vector of random variables

   $$\mathbf{x} = (x_1, \ldots, x_n),$$

   which can be written either as a row or column. These appear naturally when the outcome of an experiment are vector-valued (e.g., the experiment may be to measure the velocity of the particle, which as three components $\mathbf{v} = (v_x, v_y, v_z)$).

   (b) **Joint cumulative probability function** – The *joint cumulative probability function* of a random vector $\mathbf{x}$ is defined as

   $$F_{\mathbf{x}}(\hat{\mathbf{x}}) = \Pr(x_1 \le \hat{x}_1 \text{ and } x_2 \le \hat{x}_2 \text{ and } \cdots \text{ and } x_n \le \hat{x}_n),$$

   where $\hat{\mathbf{x}} = (\hat{x}_1, \ldots, \hat{x}_n)$.

   (c) **Joint probability density function** – The joint PDF of the random vector $\mathbf{x}$ is a function $p_{\mathbf{x}}$ defined by: They probability that, in a given trial, $\mathbf{x}(\xi)$ is in the infinitesimal interval

   $$[\hat{\mathbf{x}}, \hat{\mathbf{x}} + d\mathbf{x}] \equiv [\hat{x}_1, \hat{x}_1 + dx_1] \times \cdots \times [\hat{x}_n, \hat{x}_n + dx_n]$$

6

is

$$\Pr([\hat{\mathbf{x}}, \hat{\mathbf{x}} + d\mathbf{x}]) = f_{\mathbf{X}}(\hat{\mathbf{x}})d\mathbf{x}$$

The PDF $f_{\mathbf{X}}$ satisfies

$$\int_{\mathbb{R}^n} p_{\mathbf{x}}(\mathbf{x}')d\mathbf{x}' = \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} f_{\mathbf{X}}(x_1', \ldots, x_n')dx_1' \ldots dx_n' = 1$$

and is related to the joint CDF by

$$F_{\mathbf{x}}(\hat{\mathbf{x}}) = \int_{-\infty}^{\hat{x}_1} \cdots \int_{-\infty}^{\hat{x}_n} p_{\mathbf{x}}(x_1', \ldots, x_n')dx_1' \ldots dx_n'$$

The probability that the random vector lies in any region $\Omega \subset \mathbb{R}^n$ can be computed,

$$\Pr(\mathbf{x} \in \Omega) = \int_{\Omega} f_{\mathbf{X}}(\mathbf{x}')d\mathbf{x}'$$

(d) **Independent random variables** – The (scalar) random variables $x_1, \ldots x_n$ are said to be *independent* if

$$p_{\mathbf{X}}(\hat{\mathbf{x}}) = p_{X_1}(\hat{x}_1)p_{X_2}(\hat{x}_2)\cdots p_{X_n}(\hat{x}_n)$$

i.e. the distribution can be factored.

(e) **Mean** – The mean of a random vector $\mathbf{x}$ is

$$\langle \mathbf{x} \rangle = \int_{\mathbb{R}^n} \mathbf{x}' f_{\mathbf{X}}(\mathbf{x}')d\mathbf{x}'$$

Note that this is a compact way to represent the mean of each component of the vector

$$\langle x_i \rangle = \int_{\mathbb{R}^n} x_i' p_{\mathbf{x}}(\mathbf{x}')d\mathbf{x}'$$

(f) **Covariance** – The covariance of the two (scalar) random variables x_i and x_j is

$$C_{ij} = \langle (x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle) \rangle$$
$$= \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$$

The variance of x_i is $C_{ii} = \sigma_{x_i}^2$. The *covariance matrix* of the random vector $\mathbf{x}$ is

$$\mathbf{C} = \begin{bmatrix} C_{11} & \cdots & C_{1n} \\ \vdots & \ddots & \vdots \\ C_{1n} & \cdots & C_{nn} \end{bmatrix}$$

Note that the covariance matrix is symmetric, positive definite and invertible. The random vector $\mathbf{x}$ is *uncorrelated* if $\mathbf{C}$ is diagonal. If the elements of $\mathbf{x}$ are independent, then $\mathbf{x}$ is uncorrelated.

(g) The multivariate Gaussian distribution has the following form:

$$p_{\mathbf{X}}(\mathbf{x}) = (2\pi)^{-\frac{k}{2}}|\Sigma|^{-\frac{1}{2}}\exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^{\mathrm{T}}\Sigma^{-1}(\mathbf{x} - \mu)\right)$$

where k is the dimension of $\mathbf{x}$, Σ is the covariance and μ is the average.

(h) **Sample average** – Suppose that x is a scalar random variable with PDF p_x and f is a function of x. The *sample average* of f with N samples is

$$S_f^N = \frac{1}{N} \sum_{i=1}^{N} f(x_i),$$

where $x_1, \ldots, x_N$ are N samples from p_x (equivalently, values of the random variable x in N trials, $x_i = x(\xi_i)$).

The sample average is a function of the N samples from p_x. An equivalent interpretation is that the sample average is a function of N independent random variables, all of which have the same PDF p_x. That is, S_f^N is a function of a random vector $\mathbf{x} = (x_1, \ldots, x_N)$ with

$$p_{\mathbf{x}}(\hat{\mathbf{x}}) = p_x(\hat{x}_1) p_x(\hat{x}_2) \cdots p_x(\hat{x}_N)$$

Models and Data

1. Least-squares solution

   (a) To fit a linear model, $\mathbf{y}^P = \mathbf{A}\theta$ where $\mathbf{A}$ is the data and $\mathbf{y}^P$ is the prediction, often the sum of the squared error (deviation between the prediction and data) is minimized.

   (b) This problem is formulated as:

   $$\theta_{LS} = \underset{\theta}{arg\ max}\, \mathbf{e}^{\mathrm{T}}\mathbf{e} = \underset{\theta}{arg\ max}\, (\mathbf{y} - \mathbf{A}\theta)^{\mathrm{T}}(\mathbf{y} - \mathbf{A}\theta)$$

   (c) The solution to this optimization problem is

   $$\theta = (\mathbf{A}^{\mathrm{T}}\mathbf{A})^{-1}\mathbf{A}^{\mathrm{T}}\mathbf{y}$$

2. Maximum likelihood estimation (MLE)

   (a) The maximum likelihood estimation of the parameters of a model maximizes the probability (or likelihood) of the errors between the observed dataset and the model.

   (b) The likelihood function is sometimes written $\mathcal{L}(\theta; \mathcal{D}) = p(\mathcal{D}|\theta)$ where $\mathcal{D}$ is the dataset and θ is the vector (or scalar) parameters of the model.

   (c) To do this, we define our error to be $\epsilon = \hat{\mathbf{y}}^D - \mathbf{y}^P(\theta)$ and assume that the distribution of our errors is normal. This gives pdf:

   $$p_\epsilon(\hat{\epsilon}) = (2\pi)^{-(K/2)}|\mathbf{C}|^{-(1/2)}\exp\left(-\frac{1}{2}(\hat{\epsilon} - \langle\epsilon\rangle)^{\mathrm{T}}\mathbf{C}^{-1}(\hat{\epsilon} - \langle\epsilon\rangle)\right)$$

   (d) The MLE parameters can then be defined:

   $$\theta_{MLE} = \underset{\theta}{arg\ max}\ \mathcal{L}(\theta; \mathcal{D})$$

   (e) Given the model, the covariance structure can often be estimated before data is collected. The covariance can then be used in the design of experiments. The covariance is estimated using a two-step procedure:

i. Approximate the elements of the Hessian

$$H_{mn} = 2\sum_i \frac{1}{\sigma_i^2} \frac{\partial Y_i}{\partial P_m} \frac{\partial Y_i}{\partial P_n}$$

where Y_i is the predicted value of the observable, σ_i is the estimated uncertainty in Y_i and P_m and P_n are model parameters.

ii. The Hessian matrix is diagonalized to find the eigenvectors and eigenvalues

iii. The elements of the covariance matrix are estimated

$$C_{jk} = \sum_i \frac{1}{\lambda_i} v_{ji} v_{ki}$$

where λ_i is the i^{th} eigenvalue and v_{ji} is the i^{th} element of the j^{th} eigenvector.

iv. NOTE: if the model is nonlinear, prior estimates of the parameters will be required to calculate the partial derivatives.

(f) To perform this analysis, the covariance matrix, $\mathbf{C}$ is needed. The covariance can be estimated from the data but often there is not enough information to estimate all of the parameters. Often the covariance matrix is assumed to have a diagonal structure (i.e. the observables are assumed to be independent).

3. Bayesian parameter estimation

(a) Another estimate of the model parameters arises from a Bayesian framework which considers and estimate based on the posterior distribution of the parameters.

(b) Recall that Bayes' Theorem tells us that:

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})}$$

where again $\mathcal{D}$ is the dataset and θ are the parameters. $p(\theta)$ is referred to as the prior distribution.

(c) Solving for the full distribution of the parameters given the data can be a very difficult problem however often we are interested only in the maximum a posteriori (MAP) estimate, which is the best estimate of the parameters given the data (and our prior information about the parameters). To find this MAP estimate, we first note that the denominator is not a function of the parameters therefore

$$p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta)$$

The MAP estimate is therefore

$$\theta_{MAP} = \arg\max_\theta p(\mathcal{D}|\theta)p(\theta)$$

How do we estimate $p(\mathcal{D}|\theta)$? If we assume we have normally distributed errors, by the central limit theorem, we can write

$$p\left(<y> |x, \theta\right) = (2n)^{-\frac{1}{2}} \sigma^{-1} \exp\left(-\frac{\chi^2}{2}\right)$$

9

where

$$\chi^2 = \left(\frac{<y> - f(x,\theta)}{\sigma}\right)^2$$

and $\sigma^2 = \frac{1}{N}(<y^2> - <y>^2)$ is the sample standard deviation

(d) NOTE: mathematically there are clear ties between ML and Bayesian estimation of the parameters. The underlying difference has to do with the formulation of the problem: in MLE the parameters of the distribution are assumed to be fixed but unknown, in Bayesian the parameters are treated as random variables.

Monte Carlo Methods

1. One of the applications of the Monte Carlo Methods is to approximate integrals for high dimensions.

$$I_f = \int_\Omega f(\mathbf{x})d\mathbf{x} \tag{1}$$

e.g. In the case of 1-D,

$$I_f = \int_a^b f(x)dx \tag{2}$$

2. Crude Monte Carlo Integration (Random Sampling)

(a) Generates a vector of random variable $\mathbf{x} \in \Omega$ from a uniform PDF,

$$p_{\mathbf{x}}{}^{unif}(\mathbf{x}) = \frac{1}{m(\Omega)} \tag{3}$$

where $m(\Omega)$ is a normalization constant such that $\int_\Omega p_{\mathbf{x}}{}^{unif}(\mathbf{x})d\mathbf{x} = 1$.
e.g. In the case of 1-D,

$$p_x{}^{unif}(x) = \frac{1}{(b-a)} \tag{4}$$

(b) With the uniform PDF, we can then relate the integral I_f to the expected value of f by

$$I_f = m(\Omega)\int_\Omega f(\mathbf{x})p_{\mathbf{x}}{}^{unif}(\mathbf{x})d\mathbf{x} = m(\Omega)\langle f\rangle. \tag{5}$$

where $\langle f \rangle$ denotes average of f over the volume (or its analogue in d-dimension) Ω where $p_x{}^{\text{unif}}(x)$ is a uniform PDF

(c) We determine $\langle f \rangle$ by first evaluating f at a larger number (say N) of $\mathbf{x}$'s randomly distributed over the volume Ω and then taking the average of these values.

$$I_f \approx I_f{}^N \equiv m(\Omega)\frac{1}{N}\sum_{i=1}^{N} f(\mathbf{x}_i) \tag{6}$$

(d) Algorithm:
   i. Initialize $I_f{}^N = 0$.
   ii. For $i = 1, \ldots, N$,

     A. Compute a sample $\mathbf{x}_i$ of the random vector $\mathbf{x}$ with unifom PDF

     B. Compute the function value $f_i = f(\mathbf{x}_i)$

     C. Assign $I_f{}^N := I_f{}^N + f_i$

   iii. Return $I_f{}^N := m(\Omega)I_f{}^N/N$

(e) Note that in MATLAB,® you use the function `rand()` to sample numbers uniformly between 0 and 1.

3. The Metropolis Monte Carlo Algorithm

(a) The Metropolis algorithm uses a Markov process in order to construct a sequence of configurations, $\mathbf{R_1}$, $\mathbf{R_2}$, ... that are samples of the equilibrium (or the target) PDF $p_\mathbf{r}$. Starting from some configuration $\mathbf{R}_0$, the collection of sampled $\mathbf{R}_i$'s eventually yield a distribution that evolves to the target PDF after sampling many times. This idea is the basis for the Markov Chain Monte Carlo (MCMC) method. The Metropolis algorithm is one of the sample techniques used in MCMC.

(b) There are two conditions that need to hold for a Markov process:

   i. ergodicity: With enough steps you reach any configurations from any other configurations

   ii. detailed balance: At equilibrium, the transition from one state (e.g. $\mathbf{R}'$) to another (e.g. $\hat{\mathbf{R}}$) is as probable as the reverse case.

$$p_\mathbf{r}(\hat{\mathbf{R}})T(\hat{\mathbf{R}}, \mathbf{R}') = p_\mathbf{r}(\mathbf{R}')T(\mathbf{R}', \hat{\mathbf{R}}) \tag{7}$$

(c) If N_{eq} is the number of steps taken for the Markov process to reach equilibirium, then the sequence $R_{N_{eq}+1}, R_{N_{eq}+2}, \ldots$, is a sequence of samples from the PDF $p_\mathbf{r}$

(d) The transition PDF $T(\mathbf{R}', \hat{\mathbf{R}})$ (PDF of selection configuration $\hat{\mathbf{R}}$ as a potential move, given that the system is presently in configuration $\hat{\mathbf{R}}$) have two contributions: the selction PDF and the acceptance PDF

$$T(\mathbf{R}', \hat{\mathbf{R}}) = T^s(\mathbf{R}', \hat{\mathbf{R}})T^a(\mathbf{R}', \hat{\mathbf{R}}) \tag{8}$$

(e) We choose T^s to be symmetric such that $T^s(\mathbf{R}', \hat{\mathbf{R}}) = T^s(\hat{\mathbf{R}}, \mathbf{R}')$. We can then choose T^a that satisfies the detailed balance (7). In the Metropolis algorithm, the choice of T^a is

$$T^a = \begin{cases} 1 & \text{if } p_\mathbf{r}(\hat{\mathbf{R}}) > p_\mathbf{r}(\mathbf{R}') \\ p_\mathbf{r}(\hat{\mathbf{R}})/p_\mathbf{r}(\mathbf{R}') & \text{otherwise} \end{cases} \tag{9}$$

Stochastic Chemical Kinetics

1. Chemical master equation
The master equations does NOT describe the change in $\mathbf{X}(t)$ (which is the state of the system as a function of time), as the continuum equations would, because this vector varies stochastically. Instead, the master equation describes the *grand probability function* for a series of states $\mathbf{n}$ at a time t given an initial state $\mathbf{n_0}$ at time t_0.

$$P(\mathbf{n}, t \mid \mathbf{n}_0, t_0) = \Pr(\mathbf{X}(t) = \mathbf{n} \mid \mathbf{X}(t_0) = \mathbf{n}_0)$$

This probability density function is evolving in time which can be described by the ODE:

$$\frac{dP(\mathbf{n}, t \mid \mathbf{n}_0, t_0)}{dt} = \sum_{\mu=1}^{M} \left[P(\mathbf{n} - \boldsymbol{\nu}_\mu, t \mid \mathbf{n}_0, t_0) h_\mu(\mathbf{n} - \boldsymbol{\nu}_\mu) c_\mu - P(\mathbf{n}, t \mid \mathbf{n}_0, t_0) h_\mu(\mathbf{n}) c_\mu \right]$$

The first part of the equation describes the transition from a different state through a reaction $\boldsymbol{\nu}_\mu$ to the current state. The second part is the remnant of an equation describing the probability of staying in the same state. Thus, this encompasses all the different ways that we can get to state $\mathbf{n}$ at time t.

The CME can also be stated as

$$\frac{dP_\sigma}{dt} = \sum_{\sigma'} W_{\sigma' \to \sigma} P_{\sigma'}(t) - \sum_{\sigma'} W_{\sigma \to \sigma'} P_\sigma(t)$$

where P_σ denotes the probability that the system is in configuration σ at time t and $W_{\sigma' \to \sigma}$ denotes the rate of transition from configuration σ' to configuration σ. The CME describes the entire PDF of the state $\mathbf{X}(t)$ for every time t.

The state variables of the CME are the probabilities of every possible state of the reacting system. As a trivial example, in a system with only one species A that can create or destroy itself, the CME has one ODE for the probabilities of each of the states (i.e., there is 1 molecule of A, there are 2 molecules of A, ..., there are j molecules of A, ...). When there are multiple reacting species, every possible combinations of molecule numbers must be accounted for, and the number of ODEs in the CME can easily reach 10^{40} or more. Due to the high computational costs, most researchers need to employ a Monte Carlo approach.

2. Formalism of the stochastic problem

Consider a volume V containing N chemically reacting species, $S_1, \ldots, S_N$, and denote the number of molecules at each time t by the vector

$$\mathbf{X}(t) = (X_1(t), \ldots, X_N(t)).$$

These species can undergo M chemical reactions $R_1, \ldots, R_M$. Each R_μ reaction has an associated stoichiometry vector $\boldsymbol{\nu}_\mu$. For example, if $N = 3$, then the reaction $S_1 + S_2 \to S_3$ has the stoichiometry vector $\boldsymbol{\nu} = (-1, -1, 1)$.

If the R_μ reaction occurs at $\hat{t}$, then the state vector changes according to

$$\mathbf{X}(\hat{t} + dt) = \mathbf{X}(\hat{t} - dt) + \boldsymbol{\nu}_\mu$$

where $\mathbf{X}(\hat{t} - dt)$ and $\mathbf{X}(\hat{t} + dt)$ represent the number of molecules immediately before and after the reaction occurs, respectively.

The chemical reactions are assumed to occur stochastically, according to some probability distributions.

(a) **Counting number of possible R_μ reactions** – In order for R_μ to occur, it is necessary that one molecule of each reactant species collide with each other at some time. At any given time, it is possible that many different combinations can cause the reaction. Exactly how many distinct combinations depends on how many molecules of each species is present i.e., on $\mathbf{X}(t)$.

Denote the number of unique groups of reactants that could collide to cause R_μ by $h_\mu(\mathbf{X}(t))$. We need to be careful that we properly count the discrete number of molecules (e.g., avoid double counting).

Examples:

$$S_1 + S_2 \to S_3, \qquad h_\mu(\mathbf{X}(t)) = X_1(t)X_2(t)$$

$$S_1 + S_1 \to S_3, \qquad h_\mu(\mathbf{X}(t)) = \tfrac{1}{2}X_1(t)(X_1(t) - 1)$$

$$S_1 + S_1 + S_1 \to S_3, \qquad h_\mu(\mathbf{X}(t)) = \tfrac{1}{6}X_1(t)(X_1(t) - 1)(X_1(t) - 2)$$

$$S_1 \to S_2, \qquad h_\mu(\mathbf{X}(t)) = X_1(t)$$

(b) **Fundamental hypothesis** – Suppose that at least one complete group of R_μ reactants exists in volume V. Let $\pi_\mu(t, dt)$ denote the probability that a particular one of these groups will react in the time interval $[t, t + dt]$. The fundamental hypothesis of the stochastic approach to chemical kinetics is that, *for each reaction R_μ, there is a constant c_μ such that*

$$\pi_\mu(t, dt) = c_\mu dt$$

This assumes that the probability that R_μ will occur in the interval $[t, t + dt]$ increases linearly with dt (for small enough duration dt). In the bimolecular case, the constant c_μ can be derived from the kinetic theory of gases based on the following assumptions:

1. the positions of the molecules in V are random and uniformly distributed,
2. the velocities of the molecules in V are distributed according to the Maxwell-Boltzmann distribution.

A consequence of the fundamental hypothesis is that the probability of multiple reaction events, of any kind, in $[t, t + dt]$ can be shown to scale as $O((dt)^2)$. Therefore, in the limit as $dt \to 0$, the probability of multiple reaction events tends to zero more rapidly than dt such that the probability of single reaction events dominates. For this reason, all single reaction events in $[t, t + dt]$ can be treated as mutually exclusive, because only one reaction can occur (to first order approximation).

(c) **Probability that reaction R_μ occurs in $[t, t+dt]$** – From the fundamental hypothesis, we assume that at most one reaction can occur in $[t, t + dt]$. To compute the probability that R_μ occurs in $[t, t + dt]$, recall that there are $h_\mu(\mathbf{X}(t))$ distinct groups of reactants that could possibly react in $[t, t+dt]$, each with probability $c_\mu dt$. Since these h_μ possible reactions can be assumed to be mutually exclusive, the individual probabilities can be summed to give

$$\Pr(\text{exactly 1 } R_\mu \text{ rxn occurs in } [t, t + dt] \mid \mathbf{X}(t) = \mathbf{n}) = h_\mu(\mathbf{n})c_\mu dt$$

where $\mathbf{n}$ is a vector of N integers with the i^{th} element being the number of molecules of species S_i in the reacting system at time t.

Since at most one reaction is allowed to occur $[t, t + dt]$, the occurrences of each type of reaction $R_1, \ldots, R_M$ are mutually exclusive, which implies we can sum all the probabilities

$$\Pr(\text{exactly 1 rxn occurs in } [t, t + dt] \mid \mathbf{X}(t) = \mathbf{n}) = \sum_{\mu=1}^{M} h_\mu(\mathbf{n})c_\mu dt$$

13

(d) **The probability that no reactions occur in** $[t, t + \tau]$ – In the stochastic view of chemical kinetics, there are periods of time in which nothing happens. In order to accurately simulate this situation, some characterization is needed for *when* the next reaction will occur. Formally, we can ask: *Given that* $\mathbf{X}(t) = \mathbf{n}$, *what is the probability* $P_0(\tau, \mathbf{n})$ *that no reactions occur in the volume* V *within the time interval* $[t, t + \tau]$?

We first consider the probability $P_0(\epsilon, \mathbf{n})$ where ϵ is very small. For small enough ϵ, at most one reaction can be assumed to occur in V during $[t, t + \epsilon]$ such that

$$P_0(\epsilon, \mathbf{n}) = 1 - \Pr(\text{exactly 1 rxn occurs in } [t, t + \epsilon \mid \mathbf{X}(t) = \mathbf{n})$$

$$= 1 - \sum_{\mu=1}^{M} h_\mu(\mathbf{n}) c_\mu \epsilon$$

In order to calculate $P_0(\tau, \mathbf{n})$, we divide the interval $[t, t + \tau]$ into a large number of K intervals each of length $\epsilon = \tau/K : [t, t + \epsilon], [t + \epsilon, t + 2\epsilon], \ldots, [t + (K-1)\epsilon, t + K\epsilon]$. The probability $P_0(\tau, \mathbf{n})$ is the joint probability that reactions do not occur in every interval. We can assume the process to be a *Poisson process* such that the probability that a reaction occurs in a given interval is independent of the probability of reaction occurring in other intervals, which implies that

$$P_0(\tau, \mathbf{n}) = \underbrace{P_0(\epsilon, \mathbf{n}) P_0(\epsilon, \mathbf{n}) \cdots P_0(\epsilon, \mathbf{n})}_{K \text{ times}} = \left[1 - \sum_{\mu=1}^{M} \frac{h_\mu(\mathbf{n}) c_\mu \tau}{K} \right]^K.$$

This argument is valid for any sufficiently large K such that

$$P_0(\tau, \mathbf{n}) = \lim_{K \to \infty} \left[1 - \sum_{\mu=1}^{M} \frac{h_\mu(\mathbf{n}) c_\mu \tau}{K} \right]^K = \exp\left(-\sum_{\mu=1}^{M} h_\mu(\mathbf{n}) c_\mu \tau \right)$$

Therefore, the probability that no reaction occurs during $[t, t + \tau]$ depends on all reaction parameters $\{c_\mu\}$ and decreases exponentially with the length of the interval τ.

3. Kinetic Monte Carlo Simulation

The kinetic Monte Carlo (KMC) algorithm computes a *sample* from the PDF of $\mathbf{X}(t)$ (i.e., computes a single trajectory in time). If the KMC algorithm is run a large number of times, the the frequency of observing a state, say $\mathbf{X}(\hat{t}) = \mathbf{n}^*$, approaches the probability predicted by the CME, $P(\mathbf{n}^*, \hat{t} \mid \mathbf{n}_0, t_0)$.

(a) **Determining the next reaction time** – The KMC algorithm requires the probability that the next reaction occurs in the infinitesimal time interval $[t + \tau, t + \tau + d\tau]$, denoted by $P_{\text{next}}(\tau, \mathbf{n}) d\tau$.

The PDF of the random variable τ (at a fixed $\mathbf{n}$) is given by

$$P_{\text{next}}(\tau, \mathbf{n}) = a(\mathbf{n}) \exp\left(-a(\mathbf{n}) \tau \right)$$

where $a(\mathbf{n})$ is defined as the *total reaction propensity*

$$a(\mathbf{n}) = \sum_{\mu=1}^{M} h_\mu(\mathbf{n}) c_\mu$$

Determining the next reaction time in KMC involves sampling this PDF, which can be done by generating a random number r_1 from the uniform distribution on the interval $(0, 1]$ and computing

$$\tau = \frac{\ln(1/r_1)}{a(\mathbf{n})}$$

(b) **Determining the next reaction type** – Since the possibility of multiple reactions was excluded and the next reaction time was determined by sampling the PDF $P_{\text{next}}(\tau, \mathbf{n})$, there is guaranteed to be exactly one reaction that occurs in $[t + \tau, t + \tau + d\tau]$. There are M mutually exclusive possibilities, $R_1, \ldots, R_M$, each with probability

$$h_\mu(\mathbf{n})c_\mu d\tau, \quad \mu = 1, \ldots, M$$

The probability that reaction $R_{\mu'}$ occurs is given by the ratio

$$\frac{h_{\mu'}(\mathbf{n})c_{\mu'}}{a(\mathbf{n})}$$

For a fixed $\mathbf{n}$, this ratio is the PDF of the random variable μ', which can be sampled by selecting a number r_2 from the uniform distribution on $(0, 1]$ and choosing μ' to be the smallest integer such that

$$r_2 \leq \frac{\sum_{\mu=1}^{\mu'} h_\mu(\mathbf{n})c_\mu}{a(\mathbf{n})}$$

(c) **KMC Algorithm** – We can state the KMC algorithm as

1. Initialize: $t = t_0$ and $\mathbf{X}(t) = \mathbf{n}_0$.
2. While $t < t_f$:
   (a) Form a list of all possible rates in the system $R_1, \ldots, R_M$ and compute the total reaction propensity $a(\mathbf{X}(t))$.
   (b) Sample two random numbers, r_1 and r_2, from the uniform distribution on $(0, 1]$.
   (c) Determine the reaction time as $\tau = \frac{\ln(1/r_1)}{a(\mathbf{X}(t))}$.
   (d) Determine the reaction type as the smallest μ' such that $r_2 \leq \frac{\sum_{\mu=1}^{\mu'} h_\mu(\mathbf{n})c_\mu}{a(\mathbf{n})}$.
   (e) Carry out the selected reaction event by setting $t := t + \tau$ and $\mathbf{X}(t) := \mathbf{X}(t) + \boldsymbol{\nu}_{\mu'}$.
3. **Interpreting results**
   From the KMC algorithm, we obtain a trajectory. Depending on the system and what we want to measure, we can use one long trajectory or may need to run several trajectories. To obtain estimates of the probability of a certain state n_i, we need to compute $P(\mathbf{n} = n_i)$ from the trajectory by counting the frequency of each statein $\mathbf{n}$ in the total trajectory or set of trajectories weighted by the timesteps τ. To compute estimates of the duration of certain states or the time between each state, we iterate through the trajectories computing the time we spend in each state. For simple problems these can also be estimated from the transition probabilities.

As always, this document may have typos. Please refer to your notes for the most complete recap of the course. It's been great working with all of you. Good luck with finals!

10.34 Numerical Methods Applied to Chemical Engineering
Fall 2015