UVED Resource

Plant Growth Architecture and Production Dynamics

Preliminary Course: Applied Mathematics

Authors
V. Letort-Le Chevalier, Ecole Centrale Paris;
Other contributors
E. Peynaud, M. Jaeger, CIRAD;
H. Cournède, Ecole Centrale Paris

Contents and Objectives

Presentation

This course will briefly introduce the mathematical tools needed to understand the following chapters.

It should be considered more as a "toolbox" than as an exhaustive course on those notions.

As far as possible, we will begin each section by an introductive example, to give the reader an intuitive grasp of the notions that are defined after.

Course Objectives

The aim of this course is to enable students to:

- Get practice of binomial, geometrical and negative binomial probability laws
- Understand the principles of a set of common parametric functions
- Get notions of dynamic systems formalism
- Get notions of parameter identification and practice least square estimation

Map

In this applied mathematics sub-chapter, two different sections are presented:
- Notions in probabilities
- Parameter estimation (for dynamic systems)

The objectives of these sections are practical and oriented to the development of plant growth models and their application in confrontation with experimental data.
Notion in probabilities are introduced in order to be able to build a stochastic growth model of the plant organogenesis.
Parameter functions are more related to plant functioning, such as organ sink functions and modelling of photosynthesis.
Parameter estimation is oriented to the identification of some hidden parameters of the model.

Course content map

Probabilities

Discrete random variables

Formally, a random variable *X* is a measurable function from a probability space *(Ω, F, P)* to a measurable set *(E,ε)*.

More intuitively, let us consider chance experiments (or events) with a finite number or countable infinite number of possible outcomes.
That is, the possible values might be listed, although the list might be infinite. The sample space Ω is the set of possible outcomes of the experiment.

For example, we roll a dice and the possible outcomes are 1, 2, 3, 4, 5, and 6 corresponding to the side that turns up.

A **discrete random variable** can then simply be considered as a function whose values correspond to the possible outcomes of a given chance experiment, possibly after some transformation is applied.

For example, the outcome of the roll of one dice is a random variable *X*. And the sum of the outcomes of four independent dice rolls, $Y = X1 + X1 + X2 + X3 + X4$, is also a random variable.

A famous historical example

A dice game played an important historical role at the origins of the probability theory.

Some famous letters between Pascal and Fermat were instigated by a request for help from a French nobleman and gambler, Chevalier de Méré.
It is said that de Méré had been betting that, in four rolls of a die, at least one six would turn up. He was winning consistently and, to get more people to play, he changed the game to bet that, with a simple homothetic reasoning, in 24 rolls of two dice, a pair of sixes would turn up. However, the exact probability value is a bit more complicated than that: it is said that de Méré lost with 24 and, disappointed, considered that mathematics was wrong.

In fact:

$1 - (5/6)^4 = 0.5177... > 0.5$

$1 - (35/36)^{24} = 0.4914... < 0.5$

Dices (Photo M. Jaeger, CIRAD)

Discrete Random Function

The probability mass function *p* of a discrete random variable *X* satisfies the following properties:

1. The probability that the random variable *X* can take a specific value *j* is *p(j)*. That is:

$$\forall j \in \Omega \quad P[X=j] = p(j) = p_j \geq 0$$

2. The sum of *p(j)* over all possible values of *j* is *1*, that is

$$\Sigma_j p_j = 1$$

where *j* represents all possible values that *X* can take.

A consequence is that $0 \leq p_j \leq 1$

Expected value and variance

The **mathematical expectation** also called **expected value** of a discrete random variable *X* taking values *x1, x2, x3,...* with probabilities *p1, p2, p3,...* is defined as the weighted average of all its possible values, the weights being the respective probabilities, i.e. the infinite sum:

$E[X] = \sum_{j=1,\ldots,\infty} p_j x_j$

provided that this series absolutely converges (a necessary condition for the existence of the expected value of *X*).

Intuitively, it is the value that one would expect to get in average if the random process could be repeated an infinite number of times.

For the dice roll example, it comes:

$E[X] = \sum_{j=1,\ldots,6} 1/6 \,.\, j = 3.5$

Based on that definition, one can introduce the notion of **variance** of a random variable:

$VAR(X) = E[(X - E(X))^2] = E[X^2] - (E[X])^2$

For our case of discrete random variable, this writes, with the same notations as above:

$VAR[X] = \sum_{j=1,\ldots,\infty} p_j(x_j - m)^2 = \sum_{j=1,\ldots,\infty} p_j x_j^2 - m^2$

where *m* is the expected value of *X*.

Variance measures how far the values taken by *X* are spread around its expected value.

The **standard deviation** is the square root of the variance.

For the dice example, we have:

$VAR[X] = \sum_{j=1,\ldots,6} 1/6 \,(j - 3.5)^2 = 2.92$

A few properties

Expected value and variance

For any constant *c*, we have $E[c] = c$

The expectation operator is linear, i.e. for any random variables *X* and *Y* and any constant *c*, we have:
$E[c X + Y] = c E[X] + E[Y]$

Concerning now the variance:
$VAR[a X + b] = a^2 VAR[X]$

and
$VAR[X + Y] = VAR[X] + VAR[Y] + 2 Cov(X,Y)$
where $Cov(X,Y) = E[XY] - E[X] E[Y]$ is the covariance of *X* and *Y*

Law of Large Numbers

Let X_i, $i=1, ..., n$ be independent and identically distributed random variables, with finite expected absolute value m= E(|X1|).

Let Sn = X1 + ... + Xn.

Then, Sn/n tends toward m almost surely, i.e.:

$$P\left(\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} X_i = m\right) = 1$$

It means that if the number of trials grows to infinity, then the empirical mean of the sample converges toward the expected value of the random variable.

This law can be put in relation with the classical method of Monte-Carlo simulation that uses repeated samplings to determine the properties of some phenomena.

Bernoulli trials

A first introductive example

The most famous example of a Bernoulli process is that of successive tosses of a balanced coin. On the question did the coin land heads?, we can define *p* as a probability of success and *q* as a probability of failure.
For a fair coin, we have

$$p = q = 1/2.$$

A second example

A probabilistic view of plant development consists in considering that stems elongate by the successive addition of some basic units called phytomers.
Assume that a stem has a probability *0.8* to set in place a phytomer every month. Therefore, every month, the stem can either produce a new phytomer (success event) or be in pause (failure).
The appearance of a phytomer thus follows a Bernoulli process with the probability of success being

$$p = 0.8.$$

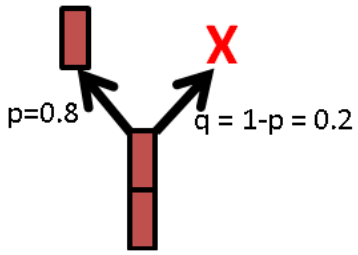

A new phytomer apparition seen as a Probability of success event (Drawing V. Letort-Lechevalier) Left scenario: the success case (with a probability 0.8 to occur); a new phytomer appears Right scenario: the failure case (with a probability 0.2 to occur); no new phytomer appears, the terminal bud rests

Definition
Repeated independent trials are called **Bernoulli trials** if there are only two possible outcomes for each trial and their probabilities remain the same throughout the trials.

The sample space of each individual trial is formed by the two possible outcomes, which are usually referred to as S (Success) and F (Failure).

The sample space of *n* Bernoulli trials contains 2^n points or successions of n symbols *S* and *F*, each representing one possible outcome of the compound experiment.
If the probability of Success is denoted as p, then the probability of failure is obviously $q = 1 - p$.

The probability of any specified sequence *{S, F, F, F, S}* is the product of the corresponding probabilities, $pqqqp = p^2(1-p)^3$.

Binomial Law

Introductive Example

Let us consider again the example of stem growth with a probability $p = 0.8$ of phytomer appearance (success) every month and observe the part of stem created after 3 months.
This part of stem can consist of 0, 1, 2, or 3 phytomers.

What are the probabilities of obtaining each of these cases?
- 0 phytomer: it means that the stem has been in pause through the 3 months. The probability is then $(1-p)^3$.

- 1 phytomer: it means that the stem has been in pause through 2 months and produced 1 phytomer during the other month. The probability of having such a sequence is $p(1-p)^2$. But several sequences can generate this observation: {Phytomer, pause, pause}, {pause, phytomer, pause} or {pause, pause, phytomer} i.e. 3 sequences. So the probability of observing only one phytomer on the stem is equal to $3p(1-p)^2$.

- 2 phytomers: it means that the stem has been in pause through one month and produced phytomers during the 2 other months. There are again 3 possible sequences that can generate this observation: the pause can be located at the first, the second or the third month. The probability is then $3(1-p)p^2$.

- 3 phytomers: it means that the stem has been producing phytomers through the 3 months. The probability is p^3.

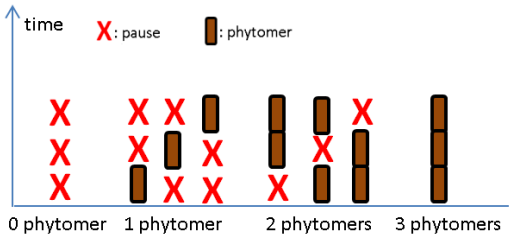

Possible stems created after 3 months (Drawing V. Letort-Chevalier, ECOLE CENTRALE PARIS)
Left: where is only one way to have no phytomers: 3 consecutive pauses
Middle: where are three ways to have a stem of 1 or 2 phytomers
Right: where is only one way to have three phytomers: 3 consecutive elongations

Definition

Frequently, we are interested in the total number of successes produced in a succession of n Bernoulli trials (but not in their order).
The number of successes can be *0, 1, 2, ..., n*.
Let *Sn* be the random variable of the number of successes in n trials.
The event -n trials result in k successes and n-k failures- can happen in as many ways as k letters S can be distributed among n places: this number is equal to the binomial coefficient

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

As we have seen before, the probability of each case is $p^k(1-p)^{n-k}$.

The probability that n Bernoulli trials with probabilities p for success and $q = 1 - p$ for failure result in k

successes and *n-k* failures is then:

$$P(S = k|n,p) = \binom{n}{k} p^k (1-p)^{n-k}$$

And this defines the binomial distribution *B(n,p)*.

Note: The Bernoulli distribution corresponds to the binomial distribution where *n = 1*.

Geometric Distribution

Consider again a sequence of Bernoulli trials $X = (X_1, X_2, \ldots)$ *with success parameter* $p \in [0,1]$.
We are interested in the random variable *N* that gives the trial number of the first success:

$N = min \{ n \in N^+, X_n = 1 \}$

$P(N = n) = p(1-p)^{n-1}$

This probability rule, applied to *n=1,..max* defines a **geometric distribution**.
The sequence of probabilities builds a **geometric sequence**.

Properties:

$E[N] = 1 / p$

$VAR[N] = (1-p) / p^2$

Negative Binomial Law

Introductive example

Let us consider again stem growth.
Assume that every month, there is a probability of 0.5 that a phytomer develops. Now, consider that, as a botanist, you are observing the stem at a given time t without knowing anything about the follow-up of the growth until that time, and want to make measurements on the last phytomer at the tip of that stem.

Problem: you do not know when that last phytomer appeared.

- It might have appeared just recently (at the current month), or one month before and the stem paused during the current month (with probability 0.5).
- It can even happen that there were two successive pauses and that the phytomer is in fact 3 months-old.
- Or be even older than that if the stem had, by (bad) chance, undergone several successive pauses.

To get information on its real chronological age, we need to assess the expected number of pauses that the stem undergoes before it creates a phytomer.

So we need to study the number of trials (months) that it takes to get one success (phytomer appearance), and more generally, *r* successes.

Consider a succession of Bernoulli trials and let us inquire how many trials it will take for the r^{th} success to turn up. It will obviously take at least *r* trials.

The probability that the r^{th} success occurs at the trial number *k* (where *k = r, r+1, ...*) equals the probability of a sequence consisting of exactly *k - r* failures and *r* successes: this occurs with probability

$(1-p)^{k-r} p^{r}$ (all the trials are assumed to be independent).

Since the r^{th} success comes last, it remains to choose the positions of *r - 1* other successes out of the remaining *k - 1* trials.

Hence the following definition

Definition

In a sequence of independent Bernoulli (*p*) trials, let the random variable *X* denote the trial at which the r^{th} success occurs, where *r* is a fixed integer.

Then *X* follows the **negative binomial distribution** *NB(r, p)* and:

$$P(X = k|\, r,p) = \binom{k-1}{r-1} p^r (1-p)^{k-r}$$

The negative binomial distribution is sometimes defined in terms of the random variable *Y* = number of failures before the r^{th} success:

$$P(Y = y|\, r, p) = \binom{r + y - 1}{y} p^r (1-p)^y,\ \ y = 0, 1, \ldots$$

This formulation is equivalent to the one above since *Y = X-r*.

Some Properties

Mean: *E(NB(r,p)) = r . (1 - p) / p*

Variance: *V(NB(r,p) = r . (1-p) / p^2*

Therefore for both Binomial and Negative Binomial law, we have: *p. E = V*

Note.
The negative binomial distribution gets its name from the alternative (and equivalent) definition:

$$P(Y = y|\, r, p) = \binom{-r}{y} p^r (-1)^y (1-p)^y$$

Dynamic systems

Useful parametric functions

A parametric function is written under the generic form *f(x,p)* where *x* is a vector of state variables and *p* is a vector of parameters.

Parametric growth functions

In the development of living organisms, measurements of growing organs have often revealed a sigmoid-like shape, with a slow initiation, a quasi-linear growth phase, and a saturation phase at the termination of expansion once physiological maturity is reached or once resources are exhausted.

Biologists are interested in quantifying the growth patterns with a few parameters.

To this end, a corpus of parametric functions, also called **growth functions**, have been developed to represent such dynamics: among others, we can cite

the Gompertz function (1825),
the Verhulst function (1838),
the Weibull function (1959),
the Richards function (1959),
the symmetrical expolinear function (Goudriaan, 1994),
or the beta density function (Johnson and Leone, 1964).

The following figure shows the shapes of these functions, used to model a plant mass evolution. The table presents the function equations.

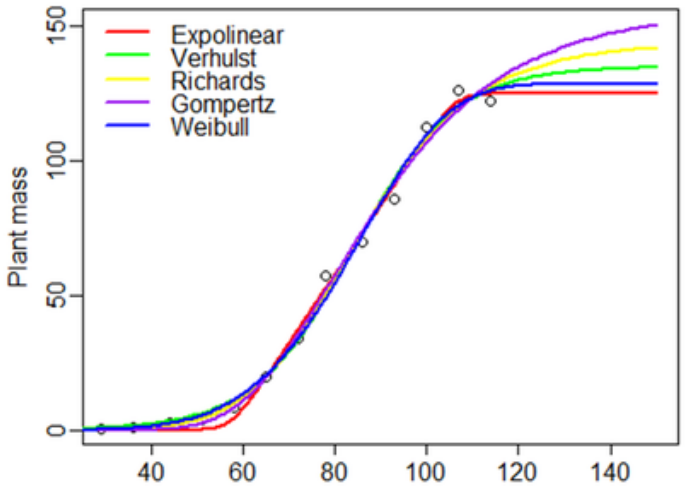

Comparison of several parametric growth functions to model plant mass evolution (dots)
Gompertz, Verhulst, Weibull, Richards, and symmetrical expolinear functions

Function	Equation
Gompertz (1825)	$w(t) = w_{max}\left[1 - e^{-e^{-r(t-t_m)}}\right]$
Verhulst (1838)	$w(t) = \frac{w_{max}}{1 + e^{-r(t-t_m)}}$
Weibull (1959)	$w(t) = w_{max}\left(1 - e^{-rt^b}\right)$
Richards (1959)	$w(t) = \frac{w_{max}}{[1 + ve^{-r(t-t_m)}]^{1/v}}$
Symmetrical expolinear (Goudriaan, 1994)	$w(t) = \frac{c}{r}\ln\frac{1 + e^{r(t-t_0)}}{1 + e^{r(t-t_0-w_{max}/c)}}$
Beta density (Johnson and Leone, 1964)	$\frac{dw}{dt} = c\left[\left(\frac{t_e - t}{t_e - t_m}\right)\left(\frac{t - t_b}{t_m - t_b}\right)^{\frac{t_m - t_b}{t_e - t_m}}\right]^d$

Parametric growth functions and their equations.

The beta density function

The beta density function has some advantages, as described by Yin (2003):

- at initial time *t = 0*, its value is zero (as for the Weibull function)
- the values of the final growth rate is null, which means that the expansion duration is finite.
- it has a high flexibility and can describe asymmetric growth trajectories (as Richards law)
- it has stable parameters for statistical estimation (similar to Gompertz or logistic functions)

Different parameterizations of the beta density function can be found.

In this course, we opt for the following one that depends only on three parameters *a, b* and *Te*, and we use a discretized form:

$$\forall n \in [0, T_e], \qquad \frac{dw}{dt}(n) = \frac{1}{N}\left(\frac{n}{T_e}\right)^{a-1}\left(1-\frac{n}{T_e}\right)^{b-1}$$

and *dw/dt(n) = 0* elsewhere.

Te represents the expansion duration while parameters *a* and *b*, usually verifying the constraint $a \geq 1$ and $b \geq 1$, drive the curve shape.

N is a multiplicative factor introduced to normalize the function so that its maximum is equal to 1:

$$N = \left(\frac{a-1}{a+b-2}\right)^{a-1}\left(\frac{b-1}{a+b-2}\right)^{b-1}$$ with $a \geq 1$ and $b \geq 1$

The following figures represent different shapes of beta density function, for several values of parameters *a* and *b*, and the shape of the integrated function (the 'growth function').

Note how asymmetrical shapes can be obtained when $a \neq b$.

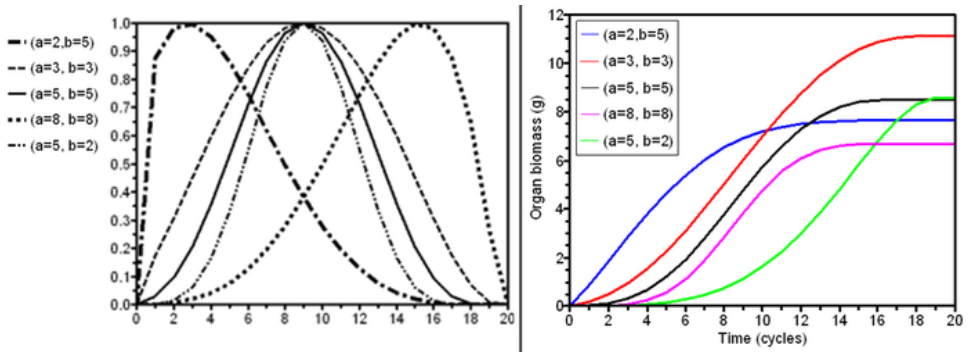

Beta density functions and their respective integrated functions

Negative exponential function

In many plant growth models, (e.g. Pilote, Ceres, STICS, GreenLab), the biomass production at time *t*, *Q(t)*, is expressed as a parametric function of the leaf area index, the LAI μ stands for a positive proportional coefficient.

$$y = Q(t) = \mu \, . \, (1 - e^{-k \, LAI(t)})$$

This expression assumes that biomass production is determined by the amount of radiation that is intercepted by the plant/field and can be derived from the Beer-Lambert law, an empirical relationship between the absorption of a monochromatic light wave in an isotropic and homogeneous medium to its concentration.
The parameter *k* drives the curve variation, μ is a real positive value standing for a biological efficiency factor.

Note that more than 95% of the incoming radiation is intercepted when $LAI \geq -ln(0.05)/k$,
i.e. when $LAI \geq \approx 3$ if $k = 1$.

Above that value, increasing the LAI has nearly no effect on the production (phenomenon of 'LAI saturation').

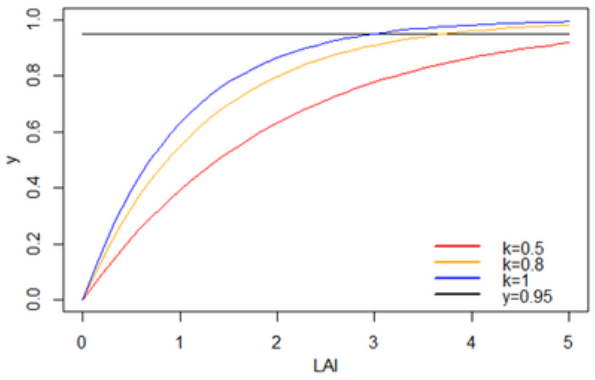

Biomass production as a function of LAI (graph V. Letort- Le Chevalier, ECOLE CENTRALE PARIS)
y stands for biomass production,
k stands for the Beer Lambert law extinction coefficient

Bibliography

Pilote: **Mailhol, J. and Gonzalez, J.** 1993. Furrow Irrigation Model for Real-Time Applications on Cracking Soils. Journal of Irrigation and Drainage Engineering, 1993, 119:5, pp. 768-783

CERES: **Jones, C.A., and J.R. Kiniry.** 1986. CERES-Maize : a simulation model of maize growth and development. Edited by C.A. Jones and J.R. Kiniry. College Station : Texas A&M University Press, 1986. 194 p.

Stics: **Brisson, N.,Gary, C., Justes, E., Roche, R., Mary, B., Ripoche, D., Zimmer, D., Sierra, J., Bertuzzi, P., Burger, P., Bussière, F. Cabidoche, Y.M., Cellier, P., Debaeke, P., Gaudillère, J.P., Hénault, C., Maraux, F., Seguin, B., Sinoquet, H.** 2003. An overview of the crop model stics, European Journal of Agronomy, Volume 18, Issues 3-4, January 2003, Pages 309-332, ISSN 1161-0301

Greenlab: **Kang, M.G., Cournède, P.H., Mathieu, A., Letort, V., & Qi, R.** 2009. A Functional-Structural Plant Model: Theories and its Applications in Agronomy. Cao, W. and White, J. and Wang, E. Crop Modeling and Decision Support, Springer, pp. 148-160, 2009, 978-3-642-01131-3

Dynamic systems

A definition of dynamic systems: a little bit of systems theory

In 'General systems research: quo vadis? ' (General Systems: Yearbook of the Society for General Systems Research, Vol.24, 1979, pp.1-9), the famous British scientist in systems theory, Brian R. Gaines, writes the following definition of a **system**:

"A system is what is distinguished as a system. (...) Systems are whatever we like to distinguish as systems (...).

What then of some of the characteristics that we do associate with the notion of a system, some form of coherence and some degree of complexity?

The Oxford English Dictionary states that: “a system is a group, set or aggregate of things, natural or artificial, forming a connected or complex whole.
[Once] you have a system, you can study it and rationalize why you made that distinction, how you can explain it, why it is a useful one.
However, none of your post-distinction rationalization is intrinsically necessary to it being a system. They are just activities that naturally follow on from making a distinction when we take note that we have done it and want to explain to ourselves, or others, why".

According to that definition, a system is characterized by our ability to discriminate between what is part of the system from what is not.

System characterization (Drawing V. Letort-Le Chevalier, ECOLE CENTRALE PARIS)

Dynamic systems are mathematical objects used to model systems whose state (or instantaneous description) changes over time.

They can be continuous or discrete, depending on whether time is considered as a continuous ($t \in \Re$) or a discrete variable ($t \in N$); stochastic or deterministic, depending on whether random effects are considered or not.

Discrete dynamic systems

Mathematical definition of discrete dynamic systems

A discrete dynamic model *M(p)* consists of two elements that can be written under the following generic form:
- A set of equations, denoted by *f*, describing the system behaviour, i.e. the transition from its current state at time *n* to the following one:

$X_{n+1} = f(n, X_n, U_n, P)$

where $X_n \in \Re^{n_x}$ represents the vector of state variables of size n_x that characterizes the system at time n, X_0 being the initial conditions;

$U_n \in \Re^{n_u}$ is the vector of control variables, of size n_u, that characterizes the external factors or stimuli that are applied to the system at time n,

$P \in \Re^{n_p}$ is the vector of the model parameters, of size n_p, i.e. a set of constants that characterizes the model behaviour.

- A set of observation functions describing the relationships between the state variables and the model output $Y \in \Re^{n_o}$ that correspond to the vector of available discrete time measured quantities:

$Y(n_s, U_{ns}, P) = g(n_s, X_{ns}(U_{ns}, n_s, P)$ for each sampling time n_s with $s = 1, ..., N_s$

where X_{ns} and U_{ns} are respectively standing for the state variable and control variable at sampling time ns

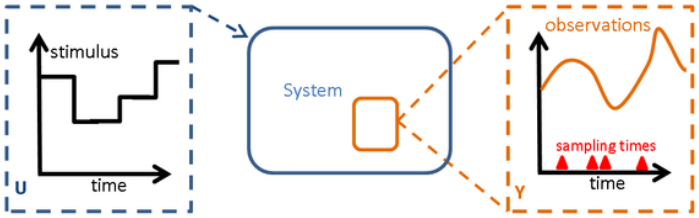

Components of a discrete dynamic system (Drawing V. Letort-Le Chevalier, ECOLE CENTRALE PARIS)

Bibliography

Gaines B.R. 1979. General systems research: quo vadis ? In: General Systems: Yearbook of the Society for General Systems Research. Vol.24, 1979, pp.1-9.

Discrete Dynamic Systems. An example

A growth model

In a plant growth model *M(P)*,

- The state variables can be for instance the amount of biomass produced at time *n*, the blade area or their nitrogen content.
- The set of parameters *P* can be considered as genetically determined.
- The environment constraints and the cultural practices (water or nutrient inputs, etc.) are included in the vector of control variables *U*.
- The output variables *Y* are some observables that can be measured on the plant such as the plant architecture or organ sizes.

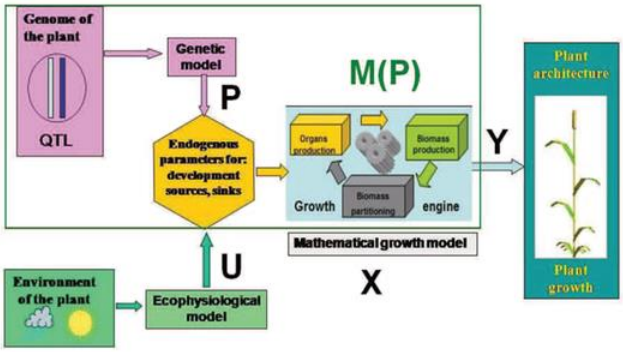

A plant growth model seen as a discrete dynamic system (Drawing V. Letort-Le Chevalier, ECOLE CENTRALE PARIS)

Equilibrium state(s)

Definition

A state X_{eq} is a controlled **equilibrium state** of a dynamic system if it exists a constant control variable U_{eq} such that:

$$X_{eq} = f\,(n, X_{eq}, U_{eq}, P\,)$$

It means that the state variables X_{eq} are constant with time under the constant stimulus U_{eq}. An equilibrium state is said to be stable if X_n stays in the neighbourhood of X_{eq} for any time *n*, if it is initially in that neighbourhood.

The number and the types of equilibrium states of a dynamic system can depend on the values of its parameter vector.

Identification

Identification of dynamic systems

Given a system *S* described by a model *M(X, P, U)* and a set of measured experimental data Y^{obs}, the **parameter estimation** of *M* consists in finding a set of parameters *P* that minimizes a distance measure between Y^{obs} and $Y = g(n, X_n(U_n, P, n), P)$.

A requested preliminary is to check the parameter identifiability.

The parameter vector *P* is structurally globally **identifiable** if for (almost*) [Note: except on a negligible subset] any P^* in the definition set of the parameter vector,

$M(P) = M(P^*) \Longrightarrow P = P^*$

An alternative and equivalent way of seeing it is to say that it should not be possible to find two different parameter vectors that would give exactly the same model outputs if simulated under the same conditions.

Studying the parameter identifiability of a model amounts to answering the question "will it be possible to uniquely estimate the model unknown parameters under given ideal (noise free and as frequent as needed) experimental conditions (stimuli and observables)? "
If the answer is "no", then the model design should be modified (e.g. some parameters should be removed) before moving to the step of parameter estimation.

Least square estimation

Parameter estimation in dynamic systems

Given a system *S* described by a model *M(X, P, U)* and a set of measured experimental data Y^{obs}, the **parameter estimation** of *M* consists in finding a set of parameters *P* that minimizes a distance measure between Y^{obs} and $Y = g(n, X_n(U_n, P, n), P)$.

A scalar function *J* needs therefore to be defined to measure the distance between model predictions and experimental data.

Quadratic cost functions are the most commonly used:

$$J(P) = (Y^{obs} - Y(P))^t \, \Omega \, (Y^{obs} - Y(P))$$

Ω is a nonnegative definite symmetric weighting matrix. The weighting coefficients located in the diagonal of the matrix are positive or zero and fixed a priori. The choice of the weights will express the relative confidence in the various experimental data and the consequent importance attached to the model performance with regard to each observable:

$\omega = 1$ assigns the same level of importance to all data

$\omega = 0$ eliminates a datum deemed not relevant

$\omega = max(Y^{obs})^2$, the square of the maximum experimental data for a given observable reduces the effect of having observations of different orders of magnitude

In the linear case, the best unbiased estimator of *P* is obtained taking weights inversely proportional to the variance of each observable.

The solution $P^* = argmin\ J(P)$ is regarded as the weighted least squares estimator.

This is an optimization problem that is generally non-linear and not analytically tractable and that can be solved using local or global numerical methods (simplex, Nelder-Mead, Newton; genetic algorithm, simulated annealing, particle swarm optimization, etc).

Supplementary resources

Inline fitting tools

Linear regression
Binomial Law fitting

Inline statistical tools

Expected value and Variance

Probability distributions and parametric functions (in English)

Keisan Online calculator. Url: http://keisan.casio.com/
Probability Functions. Url: http://keisan.casio.com/menu/system/000000000540
Logistic distribution. Url: http://keisan.casio.com/exec/system/1180573209
Beta distribution. Url: http://keisan.casio.com/exec/system/1180573225
Binomial distribution. Url: http://keisan.casio.com/exec/system/1180573199
Negative Binomial distribution. Url: http://keisan.casio.com/exec/system/1180573210
Geometric distribution. Url: http://keisan.casio.com/exec/system/1180573193
Weibull distribution. Url: http://keisan.casio.com/exec/system/1180573175

Xuru's Website. Statistical Tools. Url: http://www.xuru.org/st/PD.asp

Generic Tools. Function plotters.

MAFA Function Plotter. Url: www.mathe-fa.de/en
RechnerOnline Function Plotter. Url: http://rechneronline.de/function-graphs/
FooPlot Function Plotter. Allows Parametric. Url: http://fooplot.com/
General chart drawing. Url: http://keisan.casio.com/exec/system/1222996086

Parameter estimation

Zunzun.com, a Web fitting interface tool. Url: http://zunzun.com/
Nonlinear Least Squares Regression. Url: http://statpages.org/nonlin.html

Recommended resource

Virtual Laboratories in Probability and Statistics. Url: http://www.math.uah.edu/stat/index.html