


Basic methods in Theoretical Biology

S.A.L.M. Kooijman

Dept. Theoretical Biology
Faculty of Earth and Life Science, Vrije Universiteit, Amsterdam



This document is part of the DEB tele-course

<http://www.bio.vu.nl/thb/deb/course/>

Last update: 2011/01/07

Basic methods in Theoretical Biology discusses a basic toolkit which graduate students in Quantitative Biology, and especially in Theoretical Biology, should be able to use. The mathematical material is presented like an extended glossary; applications in biology are given in examples and exercises.

Acknowledgements I like to thank J. Ferreira, B.W. Kooi and C. Zonneveld for their helpful comments

Summary of contents:

- 1 METHODOLOGY
- 2 MATHEMATICAL TOOLKIT
- 3 MODELS FOR PROCESSES
- 4 MODEL-BASED STATISTICS

Accompanying:

EXAMPLES From biological problem via mathematics to solution

EXERCISES Motivations, given, questions, hints, answers

Contents

1	Methodology	1
1.1	Empirical cycle	1
1.2	Consistency	3
1.2.1	Dimensions	5
1.2.2	Conservation laws	7
1.3	Coherence	7
1.3.1	Scales in organization	8
1.4	Efficiency	8
1.5	Numerical behaviour	9
1.5.1	Testability	10
1.6	Experimental design	10
1.7	Identification of variables to be measured	10
1.8	Experimentation	11
1.9	Realism	11
1.9.1	Stochastic versus deterministic models	12
1.10	Logic	14
1.10.1	Propositional logic	14
1.10.2	Predicate logic	15
2	Mathematical toolkit	17
2.1	Sets	17
2.1.1	Combinatorics	18
2.2	Operators	18
2.3	Numbers	19
2.4	Functions	19
2.4.1	Trigonometric functions	20
2.4.2	Limits	21
2.4.3	Sequences and series	22
2.4.4	Differentiation	22
2.4.5	Integration	26
2.4.6	Roots	27
2.5	Matrices	28
2.5.1	Determinants	29
2.5.2	Ranks	30

2.5.3	Inverses	30
2.5.4	Eigenvalues and eigenvectors	31
2.5.5	Quadratic and bilinear forms	32
2.5.6	Vector calculus	33
2.6	Random variables and probabilities	33
2.6.1	Expectations	34
2.6.2	Examples of probability distributions	35
2.6.3	Examples of probability density functions	36
2.6.4	Conditional and marginal probabilities	37
2.6.5	Calculations with random variables	37
2.7	Numerical methods	38
2.7.1	Numerical integration	38
2.7.2	Numerical differentiation	38
2.7.3	Root finding	39
2.7.4	Extreme finding	40
3	Models for processes	41
3.1	Types of processes	41
3.1.1	Stochastic processes	41
3.2	Systems	43
3.2.1	Constraints on dynamics	44
3.2.2	Feedback	45
3.2.3	Asymptotic properties	45
4	Model-based statistics	49
4.1	Scope of statistics	49
4.2	Measurements: scales and units	49
4.3	Precision and accuracy	50
4.4	Smoothing and interpolation	53
4.5	Testing hypotheses	55
4.6	Likelihood functions	56
4.7	Large sample properties of ML estimators	57
4.8	Likelihood ratio principle	58
4.8.1	Likelihood based confidence region	59
4.8.2	Profile likelihood	60
4.9	Regression	61
4.9.1	Constant variation coefficient	61
4.10	Composite likelihoods	62
4.11	Parameter identifiability	63
4.12	Monte Carlo techniques	63
5	Notation	65
	Bibliography	67

Preface

The field of theoretical biology uses elements from methodology, mathematics, and computer science to develop new insights in biology. Such developments also require elements from biology, physics, chemistry and earth sciences. Practice learns us that this cocktail is hard to teach in a single course; it is simply too overwhelming. An extra handicap is that little knowledge of mathematics is less than adequate to deal with the complex non-linearities of life. Physics got its strength from simplification, both in theory and in experimental design. Biology, however, has little access to this powerful approach; the most simple living systems are still very complex. This is why biology still resembles a big bag of facts, semi-facts and artifacts. Yet, many theoretical biologists believe that it need not to stay like this. The purpose of this document is to present an adequate formal toolkit, that should suffice for most applications in biology.

Although many books exist on each of the topics we discuss, we found no book that just fills our educational needs. Mathematical books (and especially those on statistics) frequently have quite some material that is of little interest for the applications we have in mind, as well as that they do not present some more advanced material we think to be essential for a basic toolkit. Mathematicians use mathematics different from natural scientists, and have other purposes in mind. We here only focus on basic material for graduate students. You will find little about linear models and techniques, that dominate standard texts. The reason is that linearity hardly occurs in biology. We also omitted some standard material about computations of quantities like ranks, determinants, inverses etc, because basic computer routines are available and we do have a need for selection. You will find more on multivariate models than in elementary texts. We do realize that research frequently requires more than we offer, but the presented material should allow a rapid consultation of specialized literature.

We focus on conceptual aspects, and did not attempt to write a “stand alone” document. The serious student frequently will feel the need to consult elementary textbooks that offer more backgrounds, derivations and contexts. We suggest titles we think to be appropriate. We assume practical knowledge about Octave and/or Matlab and the availability of software package **DEBtool**, which can be downloaded from the electronic DEB laboratory at

<http://www.bio.vu.nl/thb/deb/deblab>

Design of the document on methods

A rather technical document explains elements of methodology, mathematics and computer science. The three disciplines start to blur and cross fertilize each other in chapters on modeling and statistics. The first part on methods has little material that is specific for any specialization in biology; it, therefore, remains somewhat abstract. The plan is to keep the document brief, not a collection of all-there-is, but a choice for the most basic methods, with an emphasis on concepts. The selection criterion of material is its use in the applications and exercises.

Applications

The second part of the document consists of an eventually large number of examples of application in all fields of biology. The plan is to keep each example short, starting with a biological problem and its motivation, and coming back with an answer to that problem, using pieces of the toolkit that is offered in the first part.

If new examples are included that use methods that are not discussed in the first part, the first part will be extended to include these methods.

Exercises

Besides the method document and applications data-base, an eventually substantial collection of exercises and answers is set up. Each exercise has the structure: motivation, given, question, hints, answer. The exercises can illustrate the particular method and/or an application. They can make use of public domain software.

Octave and DEBtool, which is written in Octave and Matlab, can be used to make the exercises. It is also possible to use packages like Maxima (for symbolic calculations) and AUTO for bifurcation analysis, for instance.

Use of the document

The general idea is that a graduate student, who is trained in a particular specialization in biology, can be offered a number of examples and exercises in his/her own field, together with the methods part of the document, to gain an working knowledge of Theoretical Biology.

Self improving document

We do have the idealistic view that universities have the task of optimizing the propagation of knowledge in a way that is as free of financial and cultural constraints as possible. We also do believe that collaboration leads to improvement. This is why we have setup this education project in several phases.

Phase 1: design

The executing editor first writes a first draft to structure the whole project. This is more efficient than listing the plans in detail.

Phase 2: polishing

The editorial board polishes the material, and supply additions (especially of exercises and examples).

Phase 3: self improvement

The project is now open for contributions of examples and exercises from all over the world. The editorial board will function like that of a journal and judge incoming material, seeking advice from referees. If the material requires new methods, the author should make a proposal for such a text. The first part of the document on methods will remain the responsibility of the editorial board (for the time being). If material from submissions is used in this first part, authors' name will be mentioned as contributor. The authors' name will remain associated with examples and exercises.

If the collection of examples and exercises grows large, it will be classified according to the biological subject.

Any suggestions for improvement are very welcome; please mail to deb@bio.vu.nl.

Where to get?

This document is freely down-loadable from

<http://www.bio.vu.nl/thb/course/tb/>

and will be updated and improved continuously in educational practice. We hope to stimulate the interest in the field of Theoretical Biology in this way, and to help teach a generation of students that will bring the field into blossom.

Disclaimer

This document is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

Questions?

Send to deb@bio.vu.nl.

Chapter 1

Methodology

1.1 Empirical cycle

Like it or not, but humans think in terms of models, although not everybody realizes that. The most important task of Theoretical Biology is to make implicit assumptions explicit, so that they can be replaced by others if necessary. Models have a lot in common with stories about quantities, phrased in the language of mathematics; they can have (and frequently do have) language errors, they can tell non-sense and they can be boring. They can also be exciting, however, depending on they way they are put together.

After identification of the scientific problem, the *empirical cycle* should start with the formulation of a set of assumptions, a *derivation* of a mathematical model from these assumptions, a sequence of tests on consistency, coherence, parameter sensitivity, and relevance with respect to the problem. See Figure 1.1. Most models don't need to be tested against experimental data; they simply do not pass the theoretical tests.

The second part of the empirical cycle then consists of the formulation of *auxiliary theory* for how variables in the model relate to things that can be measured, the setup of adequate experiments and/or sampling and measurement protocols to test model predictions, the collection of the measurements, and statistical tests of model predictions against measurements. These tests could reveal that the protocols have been less adequate, and should be redesigned and executed; possible inadequacies should be detected in the auxiliary theory. So inconsistencies between data and model predictions not necessarily point to inadequacies in the model itself.

If anywhere in this two-segment cycle appears the need to improve the model, it should not be changed directly, but the list of assumptions should be adapted, and the whole process should be repeated. It is a long and painstaking process, but sloppy procedures easily lead to useless results. Advocates of putting the lead of the empirical cycle in the observations, rather than in the assumptions, are frequently unaware of the implicit assumptions that need to be made to give observations a meaning. The most important aspect of modeling is to make all assumptions explicit. If modeling procedures are followed in a sloppy way, by adapting models to fit data directly, it is likely that the result will be sloppy too; one easily falls in the trap of curve-fitting. If it comes to fitting curves to data, the use of a pencil, rather than a model, is so much easier.

Empirical cycle

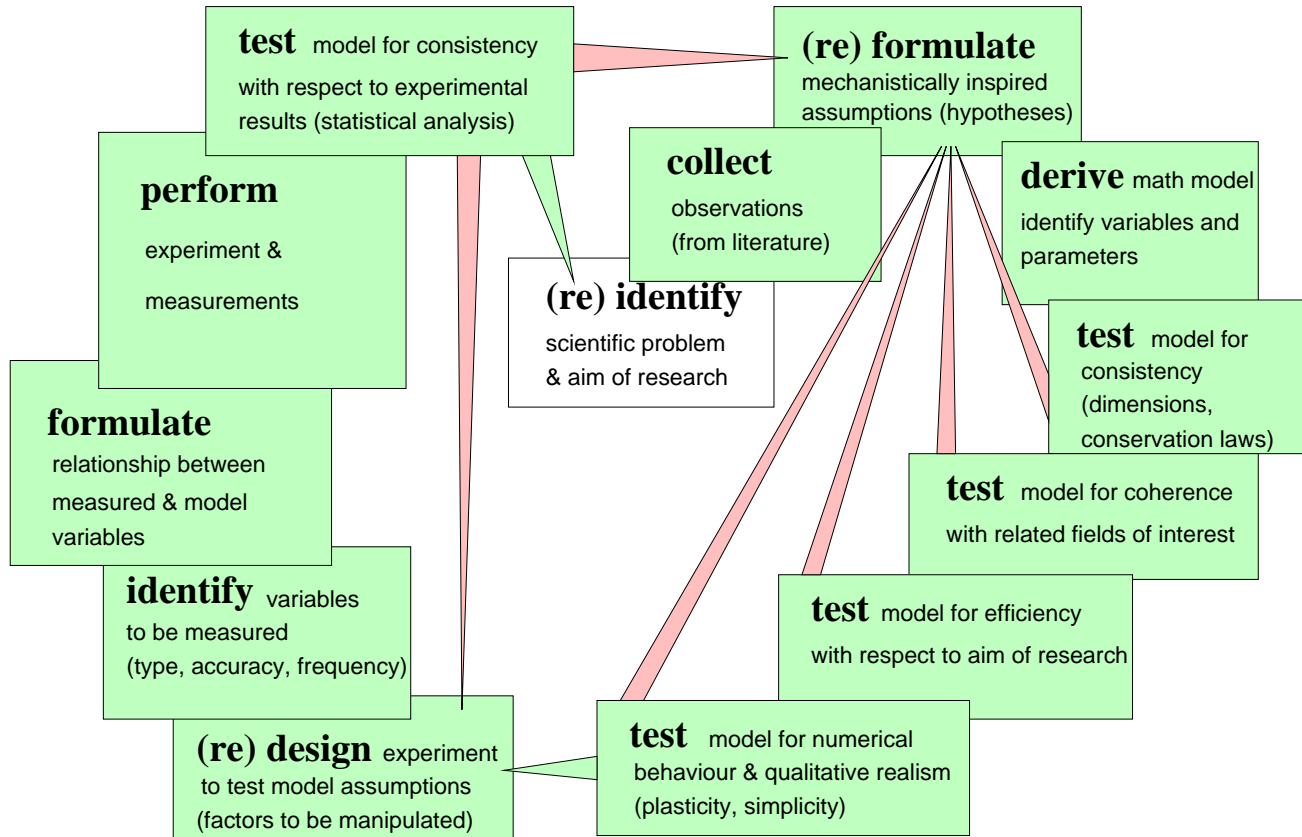


Figure 1.1: The empirical cycle as conceived by a theoretician. In the knowledge that nonsense models can easily fit any given set of data well, given enough flexibility in the parameters, realism is not the first and not the most important criterion for useful models. Lack of fit (so lack of realism) just indicates that the modeling job is not completed yet. This discrepancy between prediction and observation can be used to guide further research, which is perhaps the most useful application of models. This application to improve understanding only works if the model meets the criteria indicated in the figure; few models meet these criteria, however.

It is common practice, unfortunately, to just pose and apply a model, with little attention for the underlying assumptions. If such a model fails one of the tests, nothing is left and one should start again from scratch. There cannot be a sequence of stepwise improvements in understanding and prediction. The fact that such a model fits data is of little use, perhaps only for interpolation purposes.

Models are idealizations and, therefore, always ‘false’ in the strict sense of the word. This limits the applicability of the principle of *falsification*. A model can fit data for the wrong reasons, which means that the principle of *verification* is even more limited in applicability. This points to the criterion usefulness, to judge models, but *usefulness* is linked to a purpose. This is why a model should never be separated from its purpose. The purpose can contain elements such as increase in understanding, or in predictability. Increase in understanding can turn a useful model into a less useful one.

If a model passes all tests, including that against experimental data, there is no reason to change the assumptions, and work with them until new evidence forces reconsideration. It might seem counter intuitive, but models that fail the test against experimental data more directly serve their task in leading to greater insight, i.e. in guiding to the assumptions that require reconsideration. This obviously only works well if the steps of the formulation of assumptions has been adequate. Models are a mean in getting more insight, never an aim in themselves.

Theoretical biology specializes in the interplay between methodology, mathematics and computer science as applied in biological research. It is by its nature an interdisciplinary specialization in generalism and the natural companion of *experimental biology*. Both have complementary roles to play in the empirical cycle. We hope that Figure 1.1 makes clear that both specializations should be considered as obligate symbionts in the art of science. People with a distaste for models frequently state that ‘a model is not more than you put into it’. This is absolutely right, but instead of being a weakness, it is the single most important aspect of the use of models. Assumptions can have far reaching consequences that cannot be revealed without the use of mathematics. Put into other words: any mathematical statement is either wrong or follows from assumptions. Few people throw mathematics away for this reason. Models play an important role in the mechanism of research, as will be discussed, but also in other contexts, such as in finding answers to “what if” questions, and in solving extrapolation problems (see chapter on statistics).

The next sections highlight some steps in the empirical cycle. Table 1.1 gives some practical hints.

1.2 Consistency

Proposition X is inconsistent with proposition Y , means that they cannot both be true. Models that are internally not consistent are meaningless, so they are useless. If different assumptions are directly contradictory, inconsistency is easy to detect. In many cases, however, this is much less easy. Inconsistencies come in many forms; lack of realism (meaning: a difference between measured data and model predictions for those data) is just one form (that comes in gradations).

Table 1.1: Some practical hints for starters in science

- open a document with a unique label, your name, date, purpose; If your document is likely to contain quite some formulas, we suggest to use Latex, which is public domain
- make a list of assumptions (refer to literature items for support)
- make a list of symbols, variables *and* dimensions. Follow the mathematical rules for designing symbols; don't use names, like you will do in your computer code. Use different symbols for different dimension groups
- derive the equations, and insert any new assumptions or symbols that you need in the list. Include enough of these derivations into your document that you can understand them (much) later; check the consistency of your assumptions
- check dimensions before you proceed
- write computer code from your written formulas; we suggest to use a fourth generation language, such as Octave, which is public domain. Insert your name and date in the computer code. Refer in your code to the document where you listed the formulas. Make a link between the variables in your code and the symbols in your document
- make sure that your code is doing what your formulas prescribe. If your code is not working yet, you still don't have a problem. Problems start as soon as your code is producing something, and you have to answer the difficult question whether or not that something relates to your model
- get a numerical feel for the potential behaviour of your model by making lots of graphs using different choices of parameters
- if you don't like the numerical behaviour of your model, don't start to change the code directly. Change assumptions first, re-do your derivations, then adapt the code (including the date of creation)
- make various simplifications of your model to see what the different elements of your model do. Learn to think in terms of families of models, rather than *the* model
- plan your experiment carefully, by imaging in detail what you are going to do with the results if you would have them. Will the results answer the questions that you have?
- think of calibrating your equipment before use; does the accuracy meet your requirements? Check mass and energy balances where possible.
- specify experimental conditions (sources of materials that are used, temperature, etc); label all experimental results carefully, you might want to re-use them at a much later moment; think of using a data-base
- fit your model to the data, and make a list of parameters, estimates and units. Never insert parameter values in models, because this obscures the units.
- compare the parameter values with your expectations, based on the literature
- what is your most promising next step? Discuss your results with colleagues. Consider contacting authors of papers that you read for your work; try to be as specific as possible in the questions that you will have

\cdot	dimensionless	$\#$	number	t	time
l	length	m	mass	T	temperature

Table 1.2: Symbols for frequently used dimensions.

An example of a model inconsistency that is apparently not so easy to detect is the log-logistic model for the cumulative number of offspring per female at a standardized exposure time as a function of the concentration of test compound. This very popular model (on which much of the environmental risk assessment in the world is based) has the form $N(c) = N_0 \left(1 + (c/c_{50})^b\right)^{-1}$, where N_0 is the cumulative number of offspring in the blank, c_{50} the so-called EC50 (50% Effective Concentration) and b a parameter that relates to the slope of the concentration-response curve. The bioassay with (female) daphnids is started with a number of concentrations, and a cohort of neonates in each concentration. The individuals develop and start to reproduce after about 7 days; the bioassay runs for 21 days. The inconsistency reveals after the observation that reproduction rates tend to become constant after some time (after growth is ceased and internal concentration settled at some value). This means that the cumulative number of offspring eventually grows at a constant (concentration-dependent) rate. The implication is that, if the log-logistic model applies at some exposure time after the reproduction rates have stabilized, it cannot apply at any later exposure time. The assumption that the model applies at 21 days, together with the arbitrariness of this exposure period, in fact translates to the more stringent assumption that the model applies to all exposure times (even if it not used at other exposure times). This cannot be true, and the fact that the model fits empirical data is meaningless in the knowledge that other models fit these data too. Users of this model are probably not aware of the implicit assumptions about the reproduction process in a model that does not have time as an explicit variable. This type of problem rarely occurs if one starts from assumptions about mechanisms, rather than assuming the applicability of a model.

1.2.1 Dimensions

A dimension is an identifier for the physical nature of a quantity, see Table 1.2. A quantity of a particular dimension can be measured using several units (see Table 4.1); units determine the dimension fully. It is not necessary that all quantities in a model are measurable; the concept dimension is more general than the concept unit. Models that violate rules for dealing with dimensions are meaningless; it is a special case of inconsistency which frequently relates to errors in the translation of assumptions into a model. This does not imply that models that treat dimension well are necessarily useful models.

The elementary rules for manipulating dimensions are simple: addition and subtraction of variables are only meaningful if the dimensions of the arguments are the same, but the addition or subtraction of variables with the same dimensions is not always meaningful; meaning depends on interpretation. Multiplication and division of variables correspond with multiplication and division of dimensions. Simplifying the dimension, however, should

be done carefully. A dimension that occurs in both the numerator and the denominator in a ratio does not cancel automatically. A handy rule of thumb is that such dimensions only cancel if the sum of the variables to which they belong can play a meaningful role in the theory. The interpretation of the variable and its role in the theory always remain attached to dimensions. So the dimension of the biomass density in the environment expressed on the basis of volume is cubed length (of biomass) per cubed length (of environment); it is not dimensionless. This argument is sometimes quite subtle. The dimension of the total number of females a male butterfly meets during its lifetime is number (of females) per number (of males), as long as males and females are treated as different categories. If it is meaningful for the theory to express the number of males as a fraction of the total number of animals, the ratio becomes dimensionless.

The connection between a model and its interpretation gets lost if it contains transcendental functions of variables that are not dimensionless. Transcendental functions, such as logarithm, exponent and sinus, frequently occur in models. pH is an example, where a logarithm is taken of a variable with dimension number per cubed length ($\ln\{\#l^{-3}\}$). When it is used to specify environmental conditions, no problems arise; it just functions as a label. However, if it plays a quantitative role, we must ensure that the dimensions cancel correctly. For example, take the difference between two pH values in the same liquid. This difference is dimensionless: $\dim(\text{pH}_1 - \text{pH}_2) = \ln\{\#l^{-3}\} - \ln\{\#l^{-3}\} = \ln\{\#l^{-3}\#^{-1}l^3\} = \ln\{\cdot\} = ..$ In linear multivariate models in ecology, the pH sometimes appears together with other environmental variables, such as temperature, in a weighted sum. Here dimension rules are violated and the connection between the model and its interpretation is lost.

Another example of a model is the Arrhenius relationship, where the logarithm of a rate is linear in the inverse of the absolute temperature: $\ln \dot{k}(T) = \alpha - \beta T^{-1}$, where \dot{k} is a rate, T the absolute temperature and α and β are regression coefficients. At first sight, this model seems to violate the dimension rule for transcendental functions. However, it can also be presented as $\dot{k}(T) = \dot{k}_\infty \exp\{-T_A T^{-1}\}$, where T_A is a parameter with dimension temperature and \dot{k}_∞ is the rate at very high temperatures. In this presentation, no dimension problem arises. So, it is not always easy to decide whether a model suffers from dimension problems.

A further example of a model is the allometric function in body-size scaling relationships $\ln y(x) = \alpha + \beta \ln x$, or $y(x) = \alpha x^\beta$, where y is some variable, x has the interpretation of body weight, the parameter β is known as the scaling exponent, and α as the scaling coefficient. At first sight, this model also seems to violate the dimension rule for transcendental functions. Huxley introduced it as a solution of the differential equation $\frac{dy}{dx} = \beta \frac{y}{x}$. This equation does not suffer from dimensional problems, nor does its solution $y(x) = y(x_1) \left(\frac{x}{x_1}\right)^\beta$. However, this function has three rather than two parameters. It can be reduced to two parameters for dimensionless variables only. The crucial point is that, in most body size scaling relationships, a natural reference value x_1 does not exist for weights. The choice is arbitrary. The two-parameter allometric function violates the dimension rule for transcendental functions; uncertainty in the value of β translates into an uncertainty in the dimensions of α . Although this has been stated by many authors, the use of allometric functions is so widespread in energetics that it almost seems obligatory.

Variables are frequently transformed into dimensionless variables to simplify the model

and get rid of as many parameters as possible. This makes the structure of the model more visible, and, of course, is essential for understanding the range of possible behaviours of the model when the parameter values change. The actual values of parameters are usually known with a high degree of uncertainty and they can vary a lot. *Buckingham's theorem* states that any relationship between m variables x_i of the form $f(x_1, \dots, x_m) = 0$ can be rewritten as a relationship between $n = m - s$ dimensionless variables $y_i = h_i(x_1, \dots, x_m)$ of the form $g(y_1, \dots, y_m) = 0$, if the x 's have s different dimensions.

1.2.2 Conservation laws

Models that violate the conservation laws for mass, energy or time (or other conserved quantities) are rarely useful. It is a milder form of inconsistency. (The physical conversion between mass and energy occurs on scales in space and time that is of little relevance to life on earth.) Conservation laws can frequently be written as a constraint on state variables x_j of a system in the form $f_i(x_1, \dots, x_n) = 0$, where index i relates to the different conserved quantities (such as chemical elements, energy, etc).

Thermodynamics makes a most useful distinction between *intensive variables* – which are independent of size, such as temperature, concentration, density, pressure, viscosity, molar volume, and molar heat capacity – and *extensive variables*, which depend on size, such as mass, heat capacity and volume. Extensive variables can sometimes be added in a meaningful way if they have the same dimension, but intensive variables cannot. Concentrations, for example, can only be added when they relate to the same volume. Then they can be treated as masses, i.e. extensive variables. When the volume changes, we face the basic problem that while concentrations are the most natural choice for dealing with mechanisms, we need masses, i.e. absolute values, to make use of conservation laws. This is one of the reasons why one needs a bit of training to apply the chain rule for differentiation.

1.3 Coherence

Coherence is the natural (logical) relationship between quantities. Assumptions should not contradict 'known' relationships in the context of the model. While consistency can only be judged for rather precise quantitative propositions, coherence is weaker and is judged for more qualitative propositions. Consistency mainly applies to the assumptions in direct relationship with the problem, coherence applies to the scientific neighbourhood of the problem in a wider context. The problem that everything depends on everything else in biology has strong implications for models that represent theories. When y depends on x , it is usually not hard to formulate a set of assumptions, that imply a model that describes the relationship with acceptable accuracy. This also holds for a relationship between y and z . When more and more relationships are involved, the cumulative list of assumptions tends to grow and it becomes increasingly difficult to keep them consistent. This holds especially when the same variables occur in different relationships. Moreover, the inclusion of more variables in the model also comes with an increase in constraints that relate to

known properties of those variables.

1.3.1 Scales in organization

The field of biology ranges from molecules, via cells, individuals, population, ecosystems to system earth. These levels of organization concern scales in space as well as in time. The words of Pascal still apply:

The whole can only be understood in terms of its parts, but the parts can only be understood in the context of the whole.

Recent successes in molecular biology made holistic thinking less popular, however. Some workers seem to believe that soon they can explain all biology from the molecular level. The principle of *reduction* in science relates to the attempt to explain phenomena in terms of the smallest feasible objects. The hope for success can only be poor, however. Knowledge about technical details of engines in automobiles is extremely valuable for optimizing design, and reducing air pollution, but it is of little help to fight traffic jams. Similar relationships hold between molecular biology and ecology, these specializations focus on different space-time scales and deal with different processes that partially overlap.

Scales in space and in time are coupled in modelling because of the problem of complexity (see next section). Models with a large time scale and a small spatial scale (or vice versa) will be complex, and complex models are not very useful. Using impressive computing power, it is feasible to model water transport in the earth's oceans, which seems to defeat the coupling of scales. The modeling of this physical transport, however, involves only a limited number of parameters (and processes), given the shape of the oceans' basins, explicit external wind forcing and information on planetary rotation.

1.4 Efficiency

A model should be *well-balanced in the level of details*. It makes little sense to construct a model for x , y and z that is very detailed in the relationship between x and (y, z) , but not detailed at all in the relationships between y and z . The avoidance of such unbalance by increasing the level of detail between y and z easily leads to complex models. All details should have a necessary function in the model, both conceptually and numerically; the principle of *parsimony* is to leave out less important details. What is a detail or an essential feature depends on the problem. The *efficiency* criterion boils down to the match in essential features in the model and in the problem, which makes that the model can be used optimally to find answers to the problem.

A major trap in model building is the *complexity* caused by a large number of variables. This trap became apparent with the advent of computers, which removed the technical and practical limitations for the inclusion of many variables. Each relationship, each parameter in a relationship comes with an uncertainty, frequently an enormous one in biology. With considerable labour, it is usually possible to trim computer output to an acceptable fit with a given set of observations. This, however, gives minimal support for the realism of the

whole, which turns simulation results into a most unreliable tool for making predictions in other situations. The need for compromise between simplism and realism, makes modeling an art that is idiosyncratic to the modeler.

The only solution to the trap of complexity is to use nested *modules*. Sets of closely interacting objects are isolated from their environment and combined into a new object, a module, with simplified rules for input-output relationships. This strategy is basic to all science. A chemist does not wait for the particle physicist to finish their job, though the behaviour of the elementary particles determines the properties of atoms and molecules taken as units by the chemist. The same applies to the ecologist who does not wait for the physiologist. The existence of different specializations testifies to the relative success of the modular approach.

The problems that come with defining modules are obvious, especially when they are rather abstract. The first problem is that it is always possible to group objects in different ways to form new objects which then makes them incomparable. The problem would be easy if we could agree about the exact nature of the basic objects, but life is not that simple. The second problem with modules lies in the simplification of the input-output relationships. An approximation that works well in one circumstance can be inadequate in another. When different approximations are used for different circumstances, and this is done for several modules in a system, the behaviour of the system can easily become erratic and the approximations no longer contribute insight into the behaviour of the real thing.

In the first part of the empirical cycle, where the properties of models are analyzed, a powerful tool is to focus on the most simple models, and compare different models, where particular variables are included and excluded to study the effect of that variable. This can sometimes be done rather systematically, and families of models can be compared within a given framework. This is the happy hunting ground of *Mathematical Biology*, where model's simplicity allows the application of powerful mathematics.

1.5 Numerical behaviour

Before the realism of a model can be tested in a sensitive way, we need to study how the numerical behaviour of the model depends on the values of the variables and the parameters. Knowledge about the plasticity of the model is important in the estimation of parameter values, and in the best design of experiments. See the section on support. The rescaling of variables to dimensionless quantities is a very useful tool to reduce the complexity of the model by eliminating parameters. A very useful strategy is to choose combinations of parameter or variables values that kick out a particular mechanism, and to compare the results with other choices of values. The contribution of each mechanism to the end result can be studied this way. It frequently happens that a few combinations of a number parameters (mainly) determine the numerical behaviour, rather than each parameter separately. This reveals opportunities to simplify the model. Odd behaviour of the model can point to undesirable interactions of assumptions, but more frequently to simple programming errors. If the odd behaviour is a genuine implication of assumptions,

however, this is most helpful in the design of experiments to test its realism.

1.5.1 Testability

Models that cannot be tested against experimental results are likely to be useless. *Testability*, however, comes in gradations. In most cases assumptions can be tested indirectly only, which involves other assumptions. This complicates the process of replacement of unrealistic assumption in an attempt to find realistic ones, but this does affect the usefulness of a model.

The variables that are easy to measure or those that will be used to test the model are not always those that should be state variables. An example is metabolic rate, which is measured as the respiration rate, i.e. oxygen consumption rate or carbon dioxide production rate. The metabolic rate has different components, each of which follows simple rules. The sum of these components is then likely to behave in a less simple way in non-linear models. The same holds for, for example, dry weights, which can be decomposed into structural biomass and reserve materials. A direct consequence of such partitioning is that experimental results that only include composite variables are difficult to interpret. For mechanistic models, it is essential to use variables that are the most natural players in the game. The relationship between these variables and those to be measured is the next problem to be solved, once the model is formulated.

1.6 Experimental design

The art of experimental design fully rests on the prediction of the experimental results, and the choice of statistical procedures that will be used to evaluate the results. It is a form of reserved reasoning. The choice of experimental conditions, type of measurements to be made, details of sampling protocols to be used and other choices that have to be made can be motivated, for instance, by the minimization of the confidence intervals of particular model parameters that will be estimated from the experimental results. A problematic aspect in explicit optimization of design is that models' parameters have to be known, while the experiment is usually done because they are not known. One has to rely on guesses, which might be in the wrong ball park. Moreover, the optimization of experimental design usually also involves constraints in terms of financial costs (including effort), ethical aspects, and availability of materials. The numerical analysis of the model (see previous section) is the main source of inspiration in the design of experiments.

1.7 Identification of variables to be measured

What can be measured and the precision of measurements depend on technical possibilities and financial costs that come with their own constraints. In the most straightforward and ideal situation the variables that occur in the model can be measured directly, without interference with the system (experimental artifacts). Practice is usually remote from this ideal situation.

The usual situation is that the variables that can be measured differ from those in the model, which calls for additional modelling for how the two sets are related. These models come with new parameters, and the numerical behaviour of this model (with variables that can be measured) should again be studied to optimize the design of the experiment and reduce the complexity of the model. It frequently happens that the experiment is not a single experiment, but a set of possible very different experiments, in which different variables are measured. Some of these experiments require experimental pilot studies before the “final” experiment can be set up in an optimal way.

It is physically impossible to measure something, without interference with the system. The amount of disturbance must be evaluated in one way or another, usually by comparing results of experiments in which the disturbance is of a different nature.

Before actually performing costly (and/or time consuming) experiments it can be very useful to fake the possible measured values first, and complete the full cycle of statistical testing, using these faked values. It might be unrealistic to expect that the experiment results can possibly be satisfying, and that more effort should be invested in further optimizing the design or in the setup of alternative experiments.

1.8 Experimentation



This is not the right place to focus on the many aspects of experimentation, and e.g. procedures for calibration of measurement devices. Similar to modelling, testing mass and energy balances can be a very useful tool to check experimental results on the consistency in ways that hardly depend on modeling details.

1.9 Realism

Model predictions for measurements (or experimental results) will always differ from these measurements because models are idealisations, and repeated measurements are not identical. Whether or not a given difference is large or small depends on the specified problem. The judgement can usually be formalized in a statistical procedure for that problem, called a test, that can result in a judgement “unacceptably large”, in which case the model failed the test against experimental data. A realistic model is a model that predicts measured values with a small difference only. The measurements then support the model, and give no reason to change or replace assumptions. Such a support can, however, never prove that the assumptions are right.

The amount of *support* that a successful test of a model gives depends on the model structure and has an odd relationship with the ability to estimate parameters: the better one can estimate parameters, the less support a successful test of a model gives. This is a rather technical but vital point in work with models. This can be illustrated with a simple model that relates y to x , and which has a few parameters, to be estimated on the basis of a given set of observations $\{x_i, y_i\}$. We make a graph of the model for a given interval of the argument x , and get a set of curves if we choose the different values of the parameters between realistic boundaries. Two extremes could occur, with all possibilities in between:

- The curves have widely different shapes, together filling the whole x, y -rectangular plot. Here, one particular curve will probably match the plotted observations, determining the parameters in an accurate way, but a close match gives little support for the model; if the observations were totally different, another curve, with different parameter values, would have a close match.
- The curves all have similar shapes and are close together in the x, y -rectangular plot. If there is a close match with the observations, this gives substantial support for the model, but the parameter values are not well determined by the observations. Curves with widely different parameter values fit equally well.



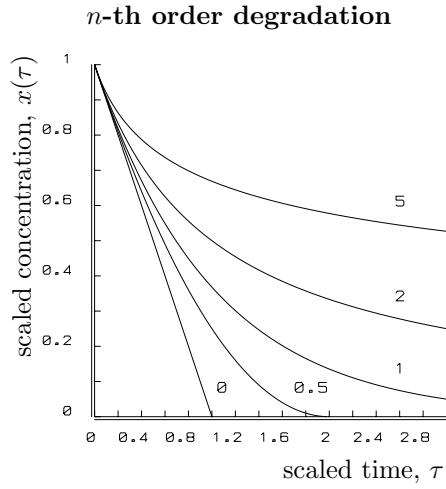
Two alternative models for biodegradation, with the same number of parameters, illustrate both situations in Figure 1.2. Of course, the choice of the model's structure is not free; it is dictated by the assumptions. So, testability is a property of the theory and nice statistical properties can combine with nasty theoretical ones and vice versa. It is essential to make this distinction.

An increase in the number of parameters usually allows models to assume a much wider range of shapes in a graph. This is closely connected with the structural property of models just mentioned. So a successful test against a set of observations gives little support for such a model, unless the set includes many variables as well. A fair comparison of models should be based on the number of parameters per variable described, not on the absolute number.

1.9.1 Stochastic versus deterministic models

Observations show scatter, which reveals itself if one variable is plotted against another. It is such an intrinsic property of biological observations that deterministic models should be considered as *incomplete*. The mechanism behind scatter is frequently the effect of a large number of factors that influence the result, but are not modeled explicitly. Think, for instance, about modeling the outcome of throwing a dice. A complex deterministic model can (in principle) predict the outcome, when the forces, the trajectory in the air, the tumbling and bouncing is modeled in great detail, including the many imperfections of dice and table. A very simple stochastic model (with the six possible outcomes having equal probability) usually works better because most parameters of the deterministic model are not known, and the process of throwing cannot be controlled in sufficient detail. This example reveals that it should usually be possible to reduce scatter (deviations between measurements and predictions by deterministic models) either by modeling more factors, or by excluding the scatter inducing factors experimentally.

Only *complete models*, i.e. those that describe observations which show scatter, can be tested. The standard way of completing deterministic models is to add 'measurement error'. The definition of a measurement error is that, if the measurements are repeated frequently



Dynamics

$$\frac{d}{dt}X = -\dot{\alpha}X^n$$

Solution

$$X(t) = \left(X_0^{1-n} - (1-n)\dot{\alpha}t\right)^{(1-n)^{-1}}$$

Special cases


$$X(t) \stackrel{n=0}{=} X_0 - \dot{\alpha}t \quad \text{for } t < X_0/\dot{\alpha}$$

$$X(t) \stackrel{n=1}{=} X_0 \exp\{-\dot{\alpha}t\}$$

Scaled solution

$$x(\tau) = (1 - (1-n)\tau)^{(1-n)^{-1}}$$

$$\text{with } x \equiv \frac{X}{X_0}; \quad \tau \equiv t\dot{\alpha}X_0^{n-1}$$



Dynamics

$$\frac{d}{dt}X = -\dot{b}\frac{X}{X_K + X}$$

Solution

$$0 = X(t) - X_0 + X_K \ln\{X(t)/X_0\} + \dot{b}t$$

Special cases

$$X(t) \stackrel{K \ll X_0}{=} X_0 - \dot{b}t \quad \text{for } t < X_0/\dot{b}$$

$$X(t) \stackrel{X_K \gg X_0}{=} X_0 \exp\{-\dot{b}t/X_K\}$$

Scaled solution

$$0 = x(\tau) - 1 + x_K \ln x(\tau) + (x_K + 1)\tau$$

$$\text{with } x \equiv \frac{X}{X_0}; \quad \tau = \frac{t\dot{b}}{X_K + X_0}; \quad x_K = \frac{X_K}{X_0}$$

Figure 1.2: The n -th order model for biodegradation of a compound X during time t is much more flexible in its morphology as a function of parameter values than the Monod model, while both models have three parameters (start concentration X_0 , a rate parameter, $\dot{\alpha}$ or \dot{b} and a shape parameter: the order n or the saturation constant X_K). Both models give identical $X(t)$ curves if $n = 0$ and $X_K \rightarrow 0$ and if $n = 1$ and $X_K \rightarrow \infty$. While all possible shapes of curves can be scaled between these two boundaries for the Monod model, many other shapes are possible for the n -th order model. This means that observations better determine the parameter values of the n -th order model, but that a good fit gives less support, compared to the Monod model. Moreover, the n -th order model suffers from dimension problems if n is not an integer, and has a more complex link with mechanisms, if any.

enough, the error will disappear in the mean of these observations. Such models are called regression models: $y_i(x_i) = f(x_i|\text{pars}) + \epsilon_i$. They are characterized by a deterministic part, here symbolized by the function f , plus a stochastic part, ϵ . The latter term is usually assumed to follow a normal probability density, with mean 0 and a fixed variance, which is one of the parameters of the model.

The interpretation of scatter as measurement error originates from physics. It is usually not realistic in biology, where many variables can be measured accurately in comparison with the amount of scatter. The observations just happen to differ from model expectations. When the scatter is large, the model is useless, despite its goodness of fit as a stochastic model. A realistic way of dealing with scatter is far from easy and usually gives rise to highly complicated models. Modelers are frequently forced to compromise between realism and mathematical over-simplicity. This further degrades the strict application of goodness of fit tests for models with unrealistic stochastic components.

1.10 Logic

Deduction is the inferring of particular instances from a general law. This is inverse to *induction*: the inference of a general law from particular instances. Most scientific reasonings have both deductive and inductive components. Deductive logic consists of two parts: propositional and predicate logic.

1.10.1 Propositional logic

Propositional logic deals with the (re)construction of statements from five elementary operations on particular statements p and q :

name	expression	symbol
negation	not- p	$\neg p$
disjunction	p or q	$p \vee q$
conjunction	p and q	$p \wedge q$
material implication	if p then q	$p \rightarrow q$
material equivalence	p if and only if q	$p \leftrightarrow q$

Based on the idea that a statement can be either “true” (T) or “false” (F), two fundamental truth tables can be constructed:

		p	q	$p \vee q$	$p \wedge q$	$p \rightarrow q$	$p \leftrightarrow q$
p	$\neg p$	T	T	T	T	T	T
T	F	T	F	T	F	F	F
F	T	F	T	T	F	T	F
F	F	F	F	F	F	T	T

The systematic application of these rules reveals, for instance, that the tautology $((p \rightarrow q) \wedge p) \rightarrow q$ is always true and the contradiction $\neg(((p \rightarrow q) \wedge p) \rightarrow q)$ is always false, for all false/true combinations for p and q .

An argument in formal logic leads from premises to a conclusion, separated by a horizontal bar. Some valid argument forms are:

modus ponens	modus tollens	disjunctive syllogism	hypothetical syllogism
$p \rightarrow q$	$p \rightarrow q$	$p \vee q$	$p \rightarrow q$
$\frac{p}{q}$	$\frac{\neg q}{\neg p}$	$\frac{\neg p}{q}$	$\frac{q \rightarrow r}{p \rightarrow r}$

Some fallacious argument forms are:

affirming the consequent	denying the antecedent	asserting an alternative
$p \rightarrow q$	$p \rightarrow q$	$p \vee q$
$\frac{q}{p}$	$\frac{\neg p}{\neg q}$	$\frac{p}{\neg q}$

1.10.2 Predicate logic

While propositional logic deals with operations on elementary statements, predicate logic deals with the internal structure of such statements. This structure is again captured in symbols.

If thing a has property P , we write the statement Pa for a one-place predicate. If relationship R exists between thing a and thing b , we write Rab for a two-place predicate. “Things” can be constants, or variables; if they are variables, statements require a quantifier. The symbol $(\forall x)$ is used for the universal quantifier, i.e. “for all x ”; so $(\forall x)(\dots x \dots)$ means “for all x , we have the statement $\dots x \dots$ ”. The symbol $(\exists x)$ represents the existential quantifier; $(\exists x)(\dots x \dots)$ means “there exists at least one x for which the statement $\dots x \dots$ holds. If a statement has no quantifier, it is called a singular statement.

Arguments related to the modus ponens are:

$$\frac{(\forall x)(Px \rightarrow Qx) \quad Pa}{Qa} \qquad \frac{(\forall x)(\forall y)((Sxy \wedge Fxy) \rightarrow Cxy) \quad Sab \wedge Fab}{Cab}$$

where x and y are variables; a and b are constants.

The following statements are pair-wise equivalent:

$$\begin{aligned} (\forall x)(Px) & \quad \neg(\exists x)(\neg Px) \\ (\exists x)(Px) & \quad \neg(\forall x)(\neg Px) \\ (\forall x)(\neg Px) & \quad \neg(\exists x)(Px) \\ (\exists x)(\neg Px) & \quad \neg(\forall x)(Px) \end{aligned}$$

Intuition will often suffice to evaluate arguments in predicate logic.

Chapter 2

Mathematical toolkit

2.1 Sets

A *set* is a collection of objects that satisfies a particular property. If such a property can be absent or present, the set can be depicted graphically in a *Venn diagram* (after John Venn, 1881), see Figure 2.1. If an object x has that property, it belongs to that set A , indicated by $x \in A$, if not it belongs to the *complement* of that set, indicated by $x \in A'$. The set of all objects U is said to be *partitioned* in the sets A and A' . The union of the two sets, $A \cup B$, is the set of objects that have at least one of two properties, intersection of the two sets, $A \cap B$, is the set of objects that have both properties. If $n(A)$ denotes the number of elements of A , we have

$$n(A \cup B) = n(A) + n(B) - n(A \cap B)$$

If all elements (also called members) of set B are also elements of A , B is said to be a *subset* of A , denoted as $B \subset A$, or $A \supset B$. If the intersection of two sets is empty, the sets are said to be *disjoint*, denoted by $A \cap B = \emptyset$.

Some rules for the three *Boolean operators* \cap , \cup and $'$ on subsets A , B and C of U are

- Idempotent laws: $A \cap A = A$; $A \cup A = A$
- Commutative laws: $A \cap B = B \cap A$; $A \cup B = B \cup A$
- Associative laws: $C \cap (A \cap B) = (C \cap A) \cap B$; $C \cup (A \cup B) = (C \cup A) \cup B$
- Absorption laws: $A \cap (A \cup B) = A$; $A \cup (A \cap B) = A$
- Modular law: If $C \supset B$, then $C \cup (A \cap B) = (C \cup A) \cap B$
- Distributive laws: $C \cap (A \cup B) = (C \cap A) \cup (C \cap B)$; $C \cup (A \cap B) = (C \cup A) \cap (C \cup B)$
- Universal bounds: $C \cap \emptyset = \emptyset$; $C \cup \emptyset = C$; $C \cap U = C$; $C \cup U = U$
- Complementarity: $C \cap C' = \emptyset$; $C \cup C' = U$




Figure 2.1: Venn's diagram depicts the set A of objects that have a certain property with a circle, the other objects don't have that property. The set B of objects that have another property can partly overlap. The union of the two sets, $A \cup B$, is the set of objects that have at least one of two properties (green and yellow), intersection of the two sets, $A \cap B$, is the set of objects that have both properties (yellow).

- Involution law: $(C')' = C$
- de Morgan's laws: $(A \cap B)' = A' \cup B'$; $(A \cup B)' = A' \cap B'$

An *interval* of real numbers is the set of all real numbers between the boundaries x_0 and x_1 . If the set includes the boundaries the interval is said to be *closed*, $[x_0, x_1]$, if not it is *open*, (x_0, x_1) . It can also be *half-open*, $(x_0, x_1]$ or $[x_0, x_1)$. An interval is said to be *unbounded* if it includes all numbers larger or smaller than some number a , and we write $[a, \infty)$, $(-\infty, a)$. If the interval includes all real numbers we write $(-\infty, \infty)$. Notice that ∞ or $-\infty$ not real numbers.

2.1.1 Combinatorics

We will only encounter very few basic results of combinatorics, the art of counting. The number of different ways to order n objects is "*n factorial*": $n! \equiv \prod_{i=1}^n i$, so e.g. $4! = 1 \times 2 \times 3 \times 4 = 24$. We have that $0! = 1! = 1$. The number of different selections of x out of n objects is "*n over x*": $\binom{n}{x} \equiv \frac{n!}{x!(n-x)!}$ for $x \leq n$, and $0! = 1$. E.g.

$$\binom{4}{2} = \frac{4!}{2!2!} = \frac{24}{2 \times 2} = 6.$$

$$\text{We have } \binom{n}{x} = \binom{n}{n-x} \text{ and } \binom{n}{x-1} + \binom{n}{x} = \binom{n+1}{x}.$$

2.2 Operators

An *operator* (or mapping or transformation) $T : \mathcal{X} \rightarrow \mathcal{Y}$ is a rule which assigns a object y in the set \mathcal{Y} (the *image*) to each objects x in the set \mathcal{X} (the *domain*). We denote the object in \mathcal{Y} assigned to x by $y = T(x)$ or by $y = Tx$. Both x and y can not only be single objects, but also lists of objects.

An operator is said to be *injective* or *one-to-one* if different elements in the domain have different images. An operator is *surjective*, or a mapping of X onto Y , if every $y \in Y$ is the image of at least on $x \in X$. An operator is said to be *bijective* if it is both injective and surjective. For such mapping, an inverse mapping exists.

2.3 Numbers

A *scalar* is a number (in contrast to e.g. a matrix, which will be discussed later). Special types of numbers are

- *Real numbers* are the familiar numbers, which can be positive, zero as well as negative. They have infinitely many decimals. If all decimals after a certain decimal are zero, the number is said to have a terminating decimal expansion. The *absolute value*, $|a|$ of real number a is a if $a \geq 0$ or $-a$ if $a < 0$. Special types of real numbers are

Integers are the numbers $0, \pm 1, \pm 2, \dots$.

Rational numbers are ratios of integers (so they include the integers).

- *Complex numbers* are numbers of the form $c = a + bi$, where real number $\text{Re}(c) = a$ is called the *real part*, real number $\text{Im}(c) = b$ is called the *imaginary part*, and i is defined as $i = \sqrt{-1}$ (so $i^2 = -1$). The complex numbers include the real ones ($b = 0$). The *conjugate* of $a + bi$ is defined to be $a - bi$; both numbers form a *conjugated pair*. Simple rules exist for operations with complex numbers

Addition: $(x + iy) + (u + iv) = x + u + i(y + v)$

Subtraction: $(x + iy) - (u + iv) = x - u + i(y - v)$

Multiplication: $(x + iy)(u + iv) = xu - yv + i(xv + yu)$

Division: $\frac{x+iy}{u+iv} = \frac{xu+yv}{u^2+v^2} + i \frac{yu-xv}{u^2+v^2}$

Complex numbers have important applications in e.g. the roots of functions, and the analysis of the behaviour of dynamic systems.

2.4 Functions

A *function* $f : \mathcal{X} \rightarrow \mathcal{Y}$ is an operator (so a rule) which assigns a unique number y in the set \mathcal{Y} (the *range*) to each number x in the set \mathcal{X} (the domain). We denote the number in \mathcal{Y} assigned to x by $y = f(x)$.

We call x the *independent variable* and y the *dependent variable*. The specification of f can contain constants, called *parameters*.

A function f is said to be *monotonously increasing* if $f(x) > f(x_1)$ for all $x > x_1$, *monotonously decreasing* if $f(x) < f(x_1)$, *monotonously non-decreasing* if $f(x) \geq f(x_1)$, *monotonously non-increasing* if $f(x) \leq f(x_1)$. A function f is said to be *even* if $f(-x) = f(x)$, or *odd* if $f(-x) = -f(x)$. A function is *homogeneous* of the n -th degree if $f(ax) = a^n f(x)$ for some arbitrary positive constant a .

Special types of functions are

- *Polynomials*: functions of the type $f(x) = \sum_{i=0}^n a_i x^i$, where n is called the *degree* of the polynomial if $a_n \neq 0$.
- *Rational functions* are ratios of two polynomials: $f(x) = p(x)/q(x)$, where $p(x)$ and $q(x)$ are polynomials.




Figure 2.2: A triangle with a right angle, and sides of length a , b and $c = \sqrt{a^2 + b^2}$ illustrates the basic trigonometric functions of angle θ . We have

$$\begin{aligned} \sin \theta &= b/c & \csc \theta &= c/b \\ \cos \theta &= a/c & \sec \theta &= c/a \\ \tan \theta &= b/a & \cot \theta &= a/b \end{aligned}$$

- *Algebraic functions*: which only include addition, multiplication, division, and taking powers.
- *Composition of two functions* g and h is the function $f(x) = g(h(x))$ for all x in the domain of the *inner function* h , such that $h(x)$ is in the domain of *outer function* g .
- *Transcendental functions*:
 - exponential function $f(x) = a^x$, where parameter a (the *base*) is real.
 - logarithmic function $f(x) = \log_a x$ if $a^{f(x)} = x$.
 - trigonometric functions $\sin(x)$, $\cos(x)$, $\tan(x)$, $\cot(x)$

Many functions in practice are only defined *implicitly* in the form $g(x, y) = 0$, where $y = f(x)$ is considered as a function of x provided that y is uniquely defined by this equation.

A function f is the *inverse function* of function g if $f(g(x)) = x$ for all values of x in the domain of g . We then also have that $g(f(x)) = x$. Examples are the pairs

- $y = \log x$ and $\exp y$ for $x > 0$
- $y = \arctan x$ and $\tan y$ for $-\pi/2 < x < \pi/2$
- $y = \arcsin x$ and $\sin y$ for $-\pi/2 < x < \pi/2$

2.4.1 Trigonometric functions

Figure 2.2 illustrates the definitions of the basic trigonometric functions *sine*, *cosecant*, *cosine*, *secant*, *tangent* and *cotangent*. The relationships between these functions follow from the observation that the sum of the three angles of a triangle amount to 2π (where π radians = 180 degrees). We have the following relationships

- $\csc \theta = \frac{1}{\sin \theta}$; $\sec \theta = \frac{1}{\cos \theta}$; $\cot \theta = \frac{1}{\tan \theta}$; $\tan \theta = \frac{\sin \theta}{\cos \theta}$.
- $\cos(-\theta) = \cos \theta$; $\sin(-\theta) = -\sin \theta$; $\cos(\theta + 2\pi) = \cos \theta$; $\sin(\theta + 2\pi) = \sin \theta$.
- $\cos^2 \theta + \sin^2 \theta = 1$; $1 + \tan^2 \theta = \sec^2 \theta$; $1 + \cot^2 \theta = \csc^2 \theta$.
- $\sin(\alpha + \beta) = \sin \alpha \cos \beta + \cos \alpha \sin \beta$; $\cos(\alpha + \beta) = \cos \alpha \cos \beta - \sin \alpha \sin \beta$
- $\cos^2 \theta = (1 + \cos 2\theta)/2$; $\sin^2 \theta = (1 - \cos 2\theta)/2$.

2.4.2 Limits

The number L is the *limit* of function $f(x)$ as x approaches x_1 provided that, given any number $\epsilon > 0$, there exists a number $\delta > 0$ such that $|f(x) - L| < \epsilon$ for all x such that $0 < |x - x_1| < \delta$. In the *right-hand limit* $\lim_{x \uparrow x_1}$, x_1 is approached from some value smaller than x_1 , in the *left-hand limit* $\lim_{x \downarrow x_1}$, x_1 is approached from some value larger than x_1 .

If a function increases without bound, we indicate this with the symbol ∞ , so $\lim_{x \uparrow 0} x^{-1} = -\infty$, $\lim_{x \downarrow 0} x^{-1} = \infty$, while $\lim_{x \rightarrow 0} x^{-1}$ does not exist. We have the rules for any positive number a : $a \pm \infty = \pm\infty$; $a \times \infty = \infty$; $\infty^a = \infty$. Undefined are: $\infty - \infty$, $0 \times \infty$, ∞^0 , $\frac{\infty}{\infty}$.

A function f is *continuous* in the neighbourhood of a point x_1 if $\lim_{x \rightarrow x_1} f(x)$ exists and equals $f(x_1)$; the left- and right-hand limits are the same, and called the two-sided limit.

Rules for limits are:

- Constant law: If $f(x) = f_1$, for constant x_1 , $\lim_{x \rightarrow x_1} f(x) = f_1$.
- Addition law: If $\lim_{x \rightarrow x_1} f(x) = f_1$ and $\lim_{x \rightarrow x_1} g(x) = g_1$, then $\lim_{x \rightarrow x_1} (f(x) + g(x)) = f_1 + g_1$.
- Product law: If $\lim_{x \rightarrow x_1} f(x) = f_1$ and $\lim_{x \rightarrow x_1} g(x) = g_1$, then $\lim_{x \rightarrow x_1} (f(x)g(x)) = f_1g_1$.
- Substitution law: If $\lim_{x \rightarrow x_1} g(x) = L$ and $\lim_{x \rightarrow L} f(x) = f(L)$, then $\lim_{x \rightarrow x_1} f(g(x)) = f(L)$.
- Reciprocal law: If $\lim_{x \rightarrow x_1} f(x) = f_1$ and $f_1 \neq 0$, then $\lim_{x \rightarrow x_1} \frac{1}{f(x)} = \frac{1}{f_1}$.
- Quotient law: If $\lim_{x \rightarrow x_1} f(x) = f_1$ and $\lim_{x \rightarrow x_1} g(x) = g_1$, while $g_1 \neq 0$ then $\lim_{x \rightarrow x_1} \frac{f(x)}{g(x)} = \frac{f_1}{g_1}$.
- Squeeze law: If $\lim_{x \rightarrow x_1} f(x) = L$ and $\lim_{x \rightarrow x_1} g(x) = L$, while $f(x) \leq h(x) \leq g(x)$ for x in the neighbourhood of x_1 , then $\lim_{x \rightarrow x_1} h(x) = L$.

The limit of $f(x, y)$ as (x, y) approaches (x_1, y_1) is L provided that, for every number $\epsilon > 0$, there exists a number $\delta > 0$, with the following property. If (x, y) is a point of the domain of f such that $0 < \sqrt{(x - x_1)^2 + (y - y_1)^2} < \delta$, then it follows that $|f(x, y) - L| < \epsilon$. If $L = f(x_1, y_1)$, and the limit exists, f is said to be continuous at the point (x_1, y_1) .

A limit that is of special interest is

$$\lim_{n \rightarrow \infty} (1 + n^{-1})^n = e \simeq 2.718281828459045$$

This number has the property that $\log_e e = 1$. If the base of the logarithm equals e , the base will be suppressed in the notation. We will write $\exp(x)$ for e^x , so $\exp(\log(x)) = x$.

The curve $y = f(x)$ has a *vertical tangent line* at the point $(x_1, f(x_1))$ provided that f is continuous at x_1 and $|\frac{d}{dx} f(x)| \rightarrow \infty$ as $x \rightarrow x_1$. The line $y = L$ is said to be a *horizontal asymptote* of the curve $y = f(x)$ if either $\lim_{x \rightarrow \infty} f(x) = L$ or $\lim_{x \rightarrow -\infty} f(x) = L$. The

function $f(x) = \frac{x}{x+1}$ has a vertical tangent line at $x = -1$ and a horizontal asymptote for $f(x) = 1$.

A function of x is said to be of *Order* x , indicated by $O(x)$ if $\lim_{x \downarrow 0} O(x) = 0$, and of *order* x , indicated by $o(x)$, if $\lim_{x \downarrow 0} o(x)/x = 0$. If a function of x is of order x , it is obviously also of Order x .

2.4.3 Sequences and series

A *sequence* is an ordered list of numbers (called *terms*), and is called an infinite sequence if the list does not end. An example the *Fibonacci sequence* $\{F_n\}$ which is defined as $F_1 = 1$, $F_2 = 1$, $F_{n+1} = F_n + F_{n-1}$ for $n \geq 2$. If new terms are defined as functions of earlier terms, like here, we speak of a *recursively defined sequence*.

The sequence $\{a_n\}$ *converges* to a real number L if $\lim_{n \rightarrow \infty} a_n = L$, which means that for any number $\epsilon > 0$ there exists an integer N such that $|a_n - L| < \epsilon$ for all $n \geq N$. A sequence is increasing if $a_i \leq a_j$ for all $i < j$, of decreasing if $a_i \geq a_j$.

A *series* is the sum of all terms in a sequence. A famous series is the *geometric series* $\sum_{n=0}^{\infty} ar^n = \frac{a}{1-r}$ for $|r| < 1$. The *harmonic series* $\sum_{n=1}^{\infty} n^{-1}$ *diverges* to infinity; it is a special case of the *Riemann zeta function*, which is defined as $\zeta(x) = \sum_{i=1}^{\infty} n^{-x}$, where x is a complex number.

The most important series is probably Newton's the *binomial series* from the 1660s:

$$1 + \alpha x + \frac{\alpha(\alpha-1)}{2!}x^2 + \frac{\alpha(\alpha-1)(\alpha-2)}{3!}x^3 + \dots = \sum_{i=0}^{\infty} \binom{\alpha}{i} x^i = (1+x)^\alpha.$$

A special case is for $x = 1$, and α equal to an integer n : $\sum_{i=0}^n \binom{n}{i} = 2^n$. Some well known series of finite sequences are $\sum_{i=1}^n i = n(n+1)/2$, $\sum_{i=1}^n i^2 = n(n+1)(2n+1)/6$, $\sum_{i=1}^n i^3 = n^2(n+1)^2/4$.

2.4.4 Differentiation

Single independent variable

The derivative of a continuous function f in the point x_1 is defined as

$$\frac{d}{dx}f(x_1) = f'(x_1) = \lim_{h \rightarrow 0} \frac{f(x_1+h) - f(x_1)}{h}$$

This limit not always exists; if it exists, the function f is said to be *differentiable* in the point x_1 . We will write dx for an infinitesimally small step in x . Notice that for $y = f(x)$, $\dim(y') = \dim(y)/\dim(x)$. The dimensionless quantity $x_1 f'(x_1)/f(x_1)$ is called the *elasticity* of f in the point x_1 . It is used in sensitivity studies.

The definition of differentiation implies rules that are listed in Table 2.1. When x has the interpretation of time, $\frac{d}{dx}f$ has the interpretation of a change in $y = f(x)$, which is a rate if y is some quantity.

$f(x)$	$f'(x)$	$f(x)$	$f'(x)$
a	0	x^a	ax^{a-1}
$g(x) + h(x)$	$g'(x) + h'(x)$	$g(x)h(x)$	$g'(x)h(x) + g(x)h'(x)$
$\frac{g(x)}{h(x)}$	$\frac{g'(x)}{h(x)} - \frac{g(x)h'(x)}{(h(x))^2}$	$g(h(x))$	$\frac{dg}{dh} \frac{dh}{dx}$
a^x	$a^x \log x$		
$\exp x$	$\exp x$	$\log x$	x^{-1}
$\cos x$	$-\sin x$	$\sin x$	$\cos x$
$\tan x$	$\sec^2 x$	$\cot x$	$-\sec^2 x$
$\sec x$	$\sec x \tan x$	$\csc x$	$-\csc x \cot x$

Table 2.1: Some frequently occurring rules for differentiation, where a is a parameter.

If a differentiable function $y = f(x)$ is specified implicitly in the form $F(x, y) = 0$, *implicit differentiation* of y with respect to x generally results in an differential equation of the form $g(x, y, y') = 0$.

Differentiation can be nested; the *second derivative* is

$$\frac{d^2}{dx^2}f(x) = f''(x_1) = \lim_{h \rightarrow 0} \frac{f'(x_1 + h) - f'(x_1)}{h}$$

When x has the interpretation of time, $y'' = \frac{d^2}{dx^2}f$ has the interpretation of a change in y' for $y = f(x)$, which is an acceleration if y is some quantity. If $y''(x)$ changes sign at $x = x_1$, $f(x)$ is said to have an *inflection point* at x_1 . We have $\dim(y'') = \dim(y) / \dim^2(x)$.

L'Hôpital's rule for the limit of the ratio of two differentiable functions f and g for some $x \neq a$ in some open interval containing a reads

$$\lim_{x \rightarrow a} \frac{f(x)}{g(x)} = \lim_{x \rightarrow a} \frac{f'(x)}{g'(x)} \quad \text{if } \lim_{x \rightarrow a} f(x) = 0 \text{ and } \lim_{x \rightarrow a} g(x) = 0 \text{ while } g'(a) \neq 0$$

The theorem of Brook Taylor (1685-1731) states that any bounded function f that has derivatives of all orders in some interval containing x_1 can be written as the polynomial

$$f(x) = \sum_{n=0}^{\infty} \frac{(x - x_1)^n}{n!} \frac{d^n}{dx^n} f(x_1) \quad (2.1)$$

For $x_1 = 0$, the Taylor series is called the *Maclaurin series*. The *Taylor series* is a generalization of the binomial series; for $f(x) = (1 + x)^\alpha$, we have

$$\frac{d^n}{dx^n} f(x) = \alpha(\alpha - 1)(\alpha - 2) \cdots (\alpha - n + 1)(1 + x)^{\alpha - n}.$$

The Taylor series for the exponential, sine and cosine functions amount to

$$\begin{aligned} \exp(x) &= 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots \\ \sin(x) &= x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \cdots \\ \cos(x) &= 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \cdots \end{aligned}$$

The *line tangent* $g(x)$ to $f(x)$ at the point $(x_1, f(x_1))$ is

$$g(x) = f(x_1) + (x - x_1) \frac{d}{dx} f(x_1).$$

This corresponds with the Taylor series, using the zero-th and first-order derivatives only. Likewise we arrive the *tangent parabola* by also including the second-order terms:

$$g(x) = f(a) + (x - x_1) f'(x_1) + (x - x_1)^2 f''(x_1)/2.$$

For x close enough to x_1 , and f sufficiently smooth, the function g can be taken as an *approximation* for function f in the neighbourhood of the point x_1 .

Several independent variables

The *partial derivatives*, with respect to x and with respect to y of a function $f(x, y)$ are two functions defined as

$$\frac{\partial}{\partial x} f(x, y) = \lim_{h \rightarrow 0} \frac{f(x+h, y) - f(x, y)}{h} \quad \text{and} \quad \frac{\partial}{\partial y} f(x, y) = \lim_{h \rightarrow 0} \frac{f(x, y+h) - f(x, y)}{h}$$

The partial derivatives of a function are sometimes organized in a vector, called the gradient, denoted by ∇f . The component of ∇f which lies in the direction of an arbitrary unit vector a is called the *directional derivative* and is given by $\nabla f^T a$.

If a differentiable function $z = f(x, y)$ is defined implicitly in the form of the equation $F(x, y, z) = 0$, *implicit partial differentiation* of z with respect to x and y gives

$$\frac{\partial z}{\partial x} = -\frac{\frac{\partial}{\partial x} F}{\frac{\partial}{\partial z} F} \quad \text{and} \quad \frac{\partial z}{\partial y} = -\frac{\frac{\partial}{\partial y} F}{\frac{\partial}{\partial z} F}$$

wherever $\frac{\partial}{\partial z} F \neq 0$.

The second order partial derivatives of f with respect to x and y are four functions $\frac{\partial^2}{\partial x^2} f(x, y)$, $\frac{\partial^2}{\partial x \partial y} f(x, y)$, $\frac{\partial^2}{\partial y \partial x} f(x, y)$, and $\frac{\partial^2}{\partial y^2} f(x, y)$, where $\frac{\partial^2}{\partial x \partial y} f(x, y) = \frac{\partial^2}{\partial y \partial x} f(x, y)$. The functions are usually organized in a symmetric matrix, called the matrix of partial derivatives, or the Hessian matrix, after Otto L. Hesse.

Suppose that $w = f(x, y)$ has continuous first-order partial derivatives and $x = g(t)$ and $y = h(t)$ are differential functions, then w is differentiable with respect to t and

$$\frac{dw}{dt} = \frac{\partial w}{\partial x} \frac{dx}{dt} + \frac{\partial w}{\partial y} \frac{dy}{dt}$$

which is known as the *chain rule*. The variable t is the independent variable, x and y are *intermediate variables* and w is the dependent variable. The chain rule generalizes to several independent variables to $\frac{\partial w}{\partial t_i} = \sum_j \frac{\partial w}{\partial x_j} \frac{\partial x_j}{\partial t_i}$.

The Taylor series for a function f of two variables in the point (x_1, y_1) that has derivatives of all orders in some interval containing (x_1, y_1) reads

$$f(x, y) = \sum_{n=0}^{\infty} \sum_{i=0}^n \frac{(x - x_1)^i (y - y_1)^{n-i}}{n!} \frac{\partial^i}{\partial x^i} \frac{\partial^{n-i}}{\partial y^{n-i}} f(x_1, y_1) \quad (2.2)$$

The *plane tangent* $g(x, y)$ to $f(x, y)$ at the point (x_1, y_1) is

$$g(x, y) = f(x_1, y_1) + (x - x_1) \frac{\partial}{\partial x} f(x_1, y_1) + (y - y_1) \frac{\partial}{\partial y} f(x_1, y_1)$$

This corresponds with the Taylor series, using the zero-th and first-order derivatives only.

Critical points

A *critical point* of a function at points in the plane region R within the boundary curve C is a point in the interior of region C (so not at the boundary curve) where the derivative equals zero or where not all partial derivatives exist. Critical points play an important role in finding extremes of functions.

Extremes

If c is in the closed interval $[a, b]$, which is in the domain of function f , then $f(c)$ is called the *minimum value* of $f(x)$ on $[a, b]$ if $f(c) \leq f(x)$ for all x in $[a, b]$, or the *maximum value* if $f(c) \geq f(x)$. If f is two times differentiable at c , and $f'(c) = 0$ while $f''(c) < 0$, $f(c)$ is a *local maximum*, and if $f'(c) = 0$ while $f''(c) > 0$, $f(c)$ is a *local minimum*. If the local minimum $f(c)$ is the smallest of all local minima in $[a, b]$, $f(c)$ is said to be the *absolute* or *global minimum*; If the local *maximum* $f(c)$ is the largest of all local maxima in $[a, b]$, $f(c)$ is said to be the *absolute* or *global maximum* in $[a, b]$.

Extremes of a function are in practice frequently found through critical points. The function $f(x) = x^3$ illustrates that it has a critical point at $x = 0$, but no extreme. The localization of critical points must, therefore, be followed by testing the properties of these points. A critical point (x_1, y_1) of function $f(x, y)$ is a point where the function has an extreme if the determinant Δ at the critical point is positive, i.e.

$$\Delta = \left(\frac{\partial^2}{\partial x^2} f(x_1, y_1) \right) \left(\frac{\partial^2}{\partial y^2} f(x_1, y_1) \right) - \left(\frac{\partial^2}{\partial x \partial y} f(x_1, y_1) \right)^2 > 0.$$

In the case of several variables, some variables can have a maximum at the critical point, others a minimum. If this occurs, we speak of a *saddle point*.

The extreme of an implicit function $y = f(x)$, given by $F(x, y) = 0$ subject to the *constraints* $g_i(x, y) = 0$ for $i = 1, 2, \dots$ can be found with the method of *Lagrange multipliers* (named after its inventor Joseph Louis Lagrange, 1736-1813). The method states that this extremes can be found by the critical points of

$$F(x, y) - \sum_i \lambda_i g_i(x, y) = 0$$

where the Lagrange multipliers λ_i are considered as extra variables. Implicit partial differentiation learns that we have to solve x, y and all λ_i 's from the system of equations

$$g_i(x, y) = 0; \quad \frac{\partial}{\partial x} F(x, y) = \sum_i \frac{\partial}{\partial x} g_i(x, y); \quad \frac{\partial}{\partial y} F(x, y) = \sum_i \frac{\partial}{\partial y} g_i(x, y).$$

The method generalizes to more variables in a straightforward way.

2.4.5 Integration

Single independent variable

An *antiderivative* of the function f is a function F such that $F'(x) = f(x)$, wherever $f(x)$ is defined. Every antiderivative G of f on an open interval has the form $G(x) = F(x) + c$ for a constant c . The collection of all antiderivatives of the function f is called the *indefinite integral* of f with respect to x , and is denoted by $\int f(x) dx = F(x) + c$, and dx is called the *integrand*. The *definite integral* of f between the boundaries x_0 and x_1 (also called lower and upper limits) is $\int_{x=x_0}^{x_1} f(x) dx = F(x_1) - F(x_0)$. If a function is continuous on $[a, b]$, it is integrable on that interval. For $F(x) = \int_{x_1}^x f(y) dy$, we have $\dim(F) = \dim(f) \dim(x)$.

Some properties of integrals are

- $\int_{x=x_0}^{x_1} f(x) dx = -\int_{x_1}^{x_0} f(x) dx$
- $\int_a^b c dx = c(b - a)$
- $\int_a^b cf(x) dx = c \int_a^b f(x) dx$
- $\int_a^b (f(x) + g(x)) dx = \int_a^b f(x) dx + \int_a^b g(x) dx$
- $\int_a^b f(x) dx = \int_a^c f(x) dx + \int_c^b f(x) dx$
- $\frac{d}{dx} (\int_a^x f(t) dt) = f(x)$. This property shows that derivation and integration are inverse processes (for one variable).
- $\int u dv = uv - \int v du$ for $du = u'(x) dx$ and $dv = v'(x) dx$; a relationship known as integration by parts.

If f is integrable on $[a, b]$, the *average value* \bar{y} of $y = f(x)$ is

$$\bar{y} = \frac{1}{b - a} \int_a^b f(x) dx$$

An integral of special interest is $\int_0^1 \frac{4 dx}{1+x^2} = \pi \simeq 3.1414592653589793$. Many frequently occurring functions are defined as integrals. Examples are

- (natural) logarithm: $\log x = \int_1^x \frac{dt}{t}$. Some properties are $\log xy = \log x + \log y$, $\log x/y = \log x - \log y$, $\log x^y = y \log x$, $\log e = 1$. The exponential function is the inverse function of the logarithm: $\exp(\log x) = x$.
- gamma function $\Gamma(x) = \int_0^\infty t^{x-1} \exp(-t) dt$. A property is $\Gamma(x) = (x - 1)!$ for $x = 1, 2, \dots$ and $\Gamma(1/2) = \sqrt{\pi}$.
- beta function $B(m, n) = \int_0^1 x^{m-1}(1-x)^{n-1} dx$, with the property $B(m, n) = \frac{\Gamma(m)\Gamma(n)}{\Gamma(m+n)}$.

If $T : R \rightarrow S$ transforms variable x in a variable u and $x = f(u)$, the integral over x translates into one over u as

$$\int_R F(x) dx = \int_S F(f(u))f(u) du$$

Leibniz's theorem for the differentiation of an integral reads

$$\frac{d}{dx} \int_{g_0(x)}^{g_1(x)} h(y, x) dy = \int_{g_0(x)}^{g_1(x)} h'(y, x) dy + h(y, g_1)g_1'(x) - h(y, g_0)g_0'(x)$$

Just as integration is inverse to differentiation, double integration is inverse to double differentiation

$$\frac{d^2}{dx^2} \left(\int_a^x \int_a^y f(t) dt dy \right) = f(x)$$

Several independent variables

The concept integration extends to several variables. Suppose that $f(x, y)$ is continuous on the rectangle $R = [a, b] \times [c, d]$. Then

$$\int_a^b \left(\int_c^d f(x, y) dy \right) dx = \int_c^d \left(\int_a^b f(x, y) dx \right) dy$$

is the integral of f over an area R . The expressions within the brackets are called *partial integrals* of f with respect to x or y .

Transformation of variables in multiple integrals uses the *Jacobian*, after Carl Jacobi (1804-1851), which is the determinant of the matrix of partial derivatives, also called the Jacobian matrix. If transformation $T : R \rightarrow S$ transforms variables x and y into variables u and v , and $x = f(u, v)$, $y = g(u, v)$, the Jacobian is defined as

$$J_T(u, v) = \begin{vmatrix} \frac{\partial}{\partial u} f(u, v) & \frac{\partial}{\partial v} f(u, v) \\ \frac{\partial}{\partial u} g(u, v) & \frac{\partial}{\partial v} g(u, v) \end{vmatrix}$$

The double integral over (x, y) translates into one over (u, v) as

$$\int \int_R F(x, y) dx dy = \int \int_S F(f(u, v), g(u, v)) J_T(u, v) du dv$$

This generalizes to more than two variables in a straightforward way; notice that this also covers univariate transformations.

2.4.6 Roots

The fundamental *theorem of the algebra* states that a polynomial of degree n has exactly n complex *roots*, i.e. values for the independent variable for which the dependent variable equals zero. Some roots may have the same value. The roots of the second-order polynomial $y(x) = a + bx + cx^2$ are

$$x_1 = \frac{-b - \sqrt{d}}{2a} \quad \text{and} \quad x_2 = \frac{-b + \sqrt{d}}{2a} \quad \text{where the determinant } d = b^2 - 4ac.$$

The roots are real and different for determinant $d > 0$, real and equal for $d = 0$ and complex valued and different for $d < 0$. For all polynomials holds that if a function has a

complex root, its conjugate is also a root. If the degree is odd, the polynomial has at least one real root.

The set of pairs (x, y) which satisfies $f(x, y) = c$, for some chosen constant c is called the *c-isocline* or the *c-contour*, which is often used in graphical presentations of the function f . Familiar geometrical objects are isoclines, such as

$$(x - x_1)^2 + (y - y_1)^2 = r^2$$

for a *circle* with radius r and center at (x_1, y_1) , or

$$\frac{x^2}{x_1^2} + \frac{y^2}{y_1^2} = 1$$

for an *ellipse* with foci at $(-\sqrt{a^2 - b^2}, 0)$ and $(\sqrt{a^2 - b^2}, 0)$ for $a > b$.

$$\frac{x^2}{x_1^2} - \frac{y^2}{y_1^2} = 1$$

for an *hyperbola*

This idea can be extended to more variables, such as the set of triplets (x, y, z) that satisfy

$$(x - x_1)^2 + (y - y_1)^2 + (z - z_1)^2 = r^2$$

and form a *sphere* of radius r and a center at (x_1, y_1, z_1) .

2.5 Matrices

A *matrix* A is a rectangular array of numbers a_{ij} , called *elements*, for which certain rules are defined. Index i refers to the row number, j to the column number; so a_{23} is the element in row 2 and column 3 of matrix A . The *size* of a matrix is the number of rows and columns; a (n, m) -matrix has n rows and m columns. The rules for matrices are

- Equality: $A = B$ if $a_{ij} = b_{ij}$ for all indices; the sizes of A and B must match.
- Addition of matrices: $A + B = C$ is a matrix with elements $c_{ij} = a_{ij} + b_{ij}$; the sizes of A and B must match. We have $A + B = B + A$, and $(A + B) + C = A + (B + C)$.
- Multiplication of matrices: $AB = C$ is a matrix with elements $c_{ij} = \sum_k a_{ik}b_{kj}$; the number of columns of A must equal the number of rows of B . We have $AB \neq BA$ (in general; the commutative law does not hold for matrices), and $(AB)C = A(BC)$, $(A + B)C = AC + BC$.
- Multiplication of a matrix and a scalar: $Ab = C$ is a matrix with elements $c_{ij} = a_{ij}b$. Notice that $Ab = bA$.
- Direct multiplication of two matrices: $A \otimes B = C$ is a matrix with elements $c_{ij} = a_{ij}b_{ij}$ for all indices; the sizes of A and B must match.

- Differentiation of a matrix: $\frac{d}{dx}A = B$ is a matrix with elements $b_{ij} = \frac{d}{dx}a_{ij}(x)$.

A consequence of these rules is the subtraction $A - B = C$, where C is a matrix with elements $c_{ij} = a_{ij} - b_{ij}$, and $A - (B - C) = A - B + C$. Division is for matrices an operation that is more complex; we first have to discuss some other concepts.

The elements a_{11}, a_{22}, \dots are called the diagonal elements of matrix A . The sum of the diagonal elements is called the *trace*.

Some special matrices exist: A *square matrix* is a matrix where the number of rows and columns are equal. A *diagonal matrix* is a square matrix of which the non-diagonal elements are all zero. The *identity matrix* I is a diagonal matrix of which all diagonal elements are one. It has the property $AI = IA = A$. The *zero matrix* $\mathbf{0}$ is a matrix of which all elements are zero. It has the properties $A\mathbf{0} = \mathbf{0}$, $\mathbf{0}A = \mathbf{0}$ and $A + \mathbf{0} = A$. An *upper triangular matrix* is a square matrix where all elements below the diagonal are zero; a *lower triangular matrix* is a square matrix where all elements above the diagonal are zero. A *vector* is a matrix with a single column; a *row vector* is a matrix with a single row.

If $AA = A$, and A is non zero, matrix A is said to be *idempotent*.

Transposition is interchanging rows and columns, and is indicated with a prime. So $B = A^T$ is a matrix with elements $b_{ij} = a_{ji}$. An implied property is $(AB)^T = B^T A^T$. A *symmetric matrix* is a square matrix with the property that $A = A^T$.

2.5.1 Determinants

The *determinant* of a square matrix A of size (n, n) is the sum

$$|A| = \sum_P (-1)^\alpha a_{1j_1} a_{2j_2} \cdots a_{nj_n} \quad (2.3)$$

where the indices j_i are all different (which means that each index occurs exactly one time), α is the number of inversions that are necessary to order the sequence j_1, j_2, \dots, j_n in the sequence $1, 2, \dots, n$, and P is the set of $n!$ different permutations of the sequence j_1, j_2, \dots, j_n . For matrix sizes (1,1), (2,2) and (3,3) this amounts to

$$\begin{aligned} |a_{11}| &= a_{11} \\ \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} &= a_{11}a_{22} - a_{12}a_{21} \\ \begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} &= a_{11}a_{22}a_{33} + a_{12}a_{23}a_{31} + a_{13}a_{21}a_{32} \\ &\quad - a_{13}a_{22}a_{31} - a_{11}a_{23}a_{32} - a_{12}a_{21}a_{33} \end{aligned}$$

Determinants are not defined for non-square matrices. Some properties of determinants are

- Transposition does not affect the determinant: $|A^T| = |A|$.
- The determinants of triangular matrices, and so of diagonal matrices, equal the product of the diagonal elements.

- The determinant of a matrix where we multiply a single row with a scalar c is c time the determinant of that matrix. Consequently we have $|cA| = c^n |A|$, if A is a (n, n) -matrix.
- The sign of the determinant reverses if we interchange two rows or two columns. Consequently the determinant must be zero of a matrix that has two equal rows or columns.
- The determinant is not affected by adding a multiple of some column to another column. The same applies to rows.
- If a row or a column has all elements equal to zero, the determinant is zero.
- We have $|AB| = |A||B|$

If $|A| = 0$, A is called *singular*, and if $|A| \neq 0$, A is called *nonsingular*.

2.5.2 Ranks

A row of a matrix is said to be *linearly dependent* of a set of other rows if it can be written as a weighted sum of the rows in that set. So if a_i denotes row i of matrix A , a_i depends linearly on rows a_j and a_k , if scalars w_j and w_k exist such that $a_i = w_j a_j + w_k a_k$. The same applies to columns. The *rank* of a (r, k) -matrix is the number of independent rows if $r \geq k$, or the number of independent columns if $r \leq k$; if $r = k$, these numbers are the same. If the rank equals $\min(r, k)$, the matrix is said to be of *full rank*. Some properties of ranks are

- $\text{rank}(A^T) = \text{rank}(A)$
- $\text{rank}(A^T A) = \text{rank}(A)$ and $\text{rank}(A A^T) = \text{rank}(A)$
- for nonsingular matrix B we have $\text{rank}(AB) = \text{rank}(A)$ and $\text{rank}(BA) = \text{rank}(A)$
- if a square matrix is full rank, it is nonsingular

2.5.3 Inverses

The *inverse* of a matrix A is a matrix A^{-1} with the property that $AA^{-1} = I$ (*right inverse*) or $A^{-1}A = I$ (*left inverse*); it is the matrix analogue of division (remember: $aa^{-1} = 1$ for scalars). Three situations can happen

- the inverse does not exist. If A is a square matrix, it is singular. The right and left inverses only exist if A is full rank.
- the inverse is unique, which only happens if A is nonsingular. We have $AA^{-1} = A^{-1}A = I$ (the left and right inverses are identical), $(AB)^{-1} = B^{-1}A^{-1}$, and $(Ab)^{-1} = b^{-1}A^{-1}$ for a non-zero scalar b .

- the inverse is not unique, which happens if A is not square. The left and right inverses differ, even their sizes.

An important application of inverses is in the solution of the so-called the *non-homogeneous linear system* $Ax = c$, where x is an unknown vector x , while vector c is known. An alternative definition of an inverse A^{-1} of matrix A is a matrix such that $x = A^{-1}c$ is a solution of $Ax = c$. It can be shown that for each A (rectangular or square, singular or nonsingular), there exist a unique matrix A^{-1} that satisfies 4 conditions: (1) $AA^{-1}A = A$, (2) $A^{-1}AA^{-1} = A^{-1}$, (3) $(AA^{-1})^T = AA^{-1}$, (4) $(A^{-1}A)^T = A^{-1}A$. Such a matrix A^{-1} is called a *generalized inverse*, or g-inverse or pseudo-inverse, of A . If A is nonsingular (so also square), the g-inverse of A is the (regular) inverse.

Some properties of inverses are

- $(A^{-1})^{-1} = A$
- $(A^T)^{-1} = (A^{-1})^T$
- $(cA)^{-1} = c^{-1}A^{-1}$, for a non-zero scalar c
- $A = AA^T(A^{-1})^T = (A^{-1})^T A^T A$
- $A^{-1} = A^{-1}(A^{-1})^T A^T = A^T(A^{-1})^T A^{-1}$
- $(A^T A)^{-1} = A^{-1}(A^{-1})^T$
- If $A = \sum_i A_i$, where $A_i^T A_j = 0$ and $A_i A_j^T = 0$ for $i \neq j$, then $A^{-1} = \sum_i A_i^{-1}$
- $A^{-1} = (A^T A)^{-1} A^T = A^T (A A^T)^{-1}$
- $A^{-1}A$, AA^{-1} , $I - A^{-1}A$, $I - AA^{-1}$ are all symmetric idempotent
- $(ABC)^{-1} = C^{-1}B^{-1}A^{-1}$ if A and C are full rank
- The inverse of a symmetric matrix is also symmetric
- The inverse of a diagonal matrix is a diagonal matrix with the reciprocals of the original elements

Several alternative definitions for generalized inverses exist, that satisfy other conditions than (3) and (4). The presented definition is also known as the Moore-Penrose inverse.

2.5.4 Eigenvalues and eigenvectors

The *eigen values*, also called *characteristic* or *latent roots* of a square (p, p) -matrix A the p solutions for λ of the determinantal equation

$$|A - \lambda I| = 0$$

The roots can be complex-valued and some roots can have the same value. Since λ is the root of polynomial of the degree p , see further under {27}. Associated with each eigen

value is a (right) *eigen vector*, which is a non-zero vector x with the property $Ax = \lambda x$, and a left eigen vector, which is a non-zero vector y with the property $y^T x = \lambda y^T$. If x is a (left or right) eigen vector ax is also an eigen vector for any non-zero scalar a , so eigen vectors are usually normalized to length 1, i.e. $x^T x = 1$. Some properties of eigen values and eigen vectors are

- The product of the eigenvalues of A is equal to $|A|$
- The sum of the eigenvalues of A is equal to the trace of A
- The eigen values of a symmetric matrix with real elements are all real.
- The eigen values of a positive definite matrix are all positive.
- A symmetric positive semidefinite matrix of size (n, n) and rank r has r positive eigen values and $n - r$ zero eigen values.
- The non-zero eigen values of AB are equal to the non-zero eigen values of BA . As a consequence, the traces of BA and AB are equal.
- The eigen values of a diagonal matrix are the diagonal elements.
- If A is symmetric, the eigen vectors x_i and x_j that correspond with the eigen values λ_i and λ_j for $i \neq j$ are orthogonal. That is to say, if the numerical values of λ_i and λ_j are identical, the corresponding eigen vectors need not be orthonormal, but can always be chosen to be orthogonal. This is because all weighted sums of these eigen vectors are again eigen vectors.
- Every real symmetric matrix A has an orthogonal matrix P such that $P^T A P = D$ while $P^T P = I$, where D is a diagonal matrix with eigen values on the diagonal. So, if $x = Py$, the quadratic form $x^T A x = y^T P^T A P y = y^T D y = \sum_i \lambda_i y_i^2$, where λ 's are the eigen values.

2.5.5 Quadratic and bilinear forms

The scalar function $f(x, y) = x^T A y$ is known as a *bilinear form*, and the special case $f(x) = x^T A x$ is known as *quadratic form*, for a symmetric matrix A .

A symmetric matrix A is *positive definite* if $x^T A x > 0$ for all non-zero x , or *positive semidefinite* if $x^T A x \geq 0$. The concepts *negative definite*, *negative semidefinite* are defined in a similar way, while a matrix is *indefinite* if its quadratic form can be positive as well as negative. If A is positive definite, it is also full rank.

We have that $\frac{d}{dx} x^T A x = 2Ax$.

2.5.6 Vector calculus

The *inner product* of vectors a and b of size n is defined as $a^T b$ (which is a scalar). The inner product $a^T a$ is known as the squared *length* (also called *norm*) of vector a . For size 2, this amounts to Pythagoras' theorem

$$\|y - x\| = (x^T y)^{1/2} = \left((y_1 - x_1)^2 + (y_2 - x_2)^2 \right)^{1/2}$$

The *outer product* as ab^T (which is a (n, n) -matrix)

If a vector is thought to represent the coordinates of the end point of a line segment that starts in the origin in the n -dimensional space, the cosine of the angle between vector x and y is given by

$$\cos \theta = \frac{x^T y}{(x^T x)^{1/2} (y^T y)^{1/2}}$$

If $\cos \theta = 0$, vectors x and y are said to be *orthogonal*, and if $x^T x = y^T y = 1$ as well, they are *orthonormal*. If all columns of matrix X represent orthonormal vectors, we have $XX^T = X^T X = I$, and $X^T = X^{-1}$.

For the *rotation matrix*

$$T = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix},$$

the product $y = Tx$ represents a rotation of x through an angle θ .

2.6 Random variables and probabilities

A *random variable* is a variable the values of which are given by chance. The set of possible values is called the *domain*. If the domain is a denumerable subset of the real numbers, such as $\{1, 2, 3, 4, 5, 6\}$ or $\{0, 1, 2, \dots\}$, the random variable is called *discrete* and is dimensionless. If the domain is an interval within the real numbers, like the interval $[0, 1]$ or $(0, \infty)$, the random variable is called *continuous* and can have dimensions.

Observations (or measurements) represent realizations of a random variable; a set of n similar observations is called a *sample* of size n . The *relative frequency distribution* of values in the sample converges to the probability distribution for increasing sample size.

If \underline{x} is a discrete random variable, then every element x of the domain W of \underline{x} has a certain probability to be realized, symbolized by $\Pr\{\underline{x} = x\}$. The set of all probabilities is called a *probability distribution*. Notice that $\Pr\{\underline{x} = x\} \geq 0$ and $\sum_{x \in W} \Pr\{\underline{x} = x\} = 1$.

For continuous random variables, probabilities are specified in a different way. Now for every $x \in W$ we must have $\Pr\{\underline{x} = x\} = 0$, because W is an uncountable set. The *probability density function* (p.d.f.) $f_{\underline{x}}$ specifies the probabilities in the form $\Pr\{x_0 \leq \underline{x} \leq x_1\} = \int_{x_0}^{x_1} f_{\underline{x}}(x) dx$. If the domain of a continuous \underline{x} is W , we have

$$f_{\underline{x}}(x) \geq 0 \quad \text{and} \quad \int_W f_{\underline{x}}(x) dx = 1.$$

For a p.d.f. we have $\dim(f_{\underline{x}}(x)) = 1/\dim(x)$, while probabilities are dimensionless. The product $f_{\underline{x}}(x) dx$ has the interpretation of an infinitesimally small probability $\Pr\{x \leq \underline{x} \leq x + dx\}$.

The (*cumulative*) *distribution function* of a random variable is defined as $F_{\underline{x}}(x) = \Pr\{\underline{x} \leq x\}$, and the *survivor function* of a random variable is defined as $S_{\underline{x}}(x) = \Pr\{\underline{x} > x\}$. So we have $F_{\underline{x}}(x) = 1 - S_{\underline{x}}(x)$ for all values of x . The survivor function takes its name from a very special random variable: the lifespan of an object. The survivor function then specifies the probability that the object will live longer than the value of its argument indicates. Notice that $F_{\underline{x}}(x) = \int_0^x f_{\underline{x}}(s) ds$ for a continuous random variable that takes values on the interval $[0, \infty)$, and $F_{\underline{x}}(x) = \sum_{s \leq x} \Pr\{\underline{x} = s\}$ for a discrete random variable. Distribution and survivor functions are probabilities, and therefore dimensionless, with values on the interval $[0, 1]$. They have the same meaning for discrete as for continuous variables.

2.6.1 Expectations

The *expectation* of a function of a random variable is defined as $\mathcal{E}g(\underline{x}) = \sum_x g(x) \Pr\{\underline{x} = x\}$ for discrete random variables, or $\mathcal{E}g(x) = \int_x g(x) f_{\underline{x}}(x) dx$ for continuous random variables, where the summation or integration is across all values in the domain of the random variable. The *mean* is a special case of an expectation of a function of a random variable, namely $g(x) = x$. The *variance* is another special case, namely $g(x) = (x - \mathcal{E}(x))^2$, and it is indicated with $\text{var } \underline{x}$. Notice that $\text{var } \underline{x} = \mathcal{E}\underline{x}^2 - (\mathcal{E}\underline{x})^2$, and that $\dim(\mathcal{E}(x)) = \dim(x)$ and $\dim(\text{var } \underline{x}) = \dim(x)^2$. Expectations of the function $g(x) = x^m$ are called the *m-th moment*, and of $g(x) = (x - \mathcal{E}x)^m$, the *m-th central moment*. The *standard deviation* is defined as the square root of the variance, so $\text{sd } \underline{x} = \sqrt{\text{var } \underline{x}}$; the *variation coefficient* as the ratio of the standard deviation and the mean, so $\text{vc } \underline{x} = \text{sd } \underline{x} / \mathcal{E}\underline{x}$. Notice that the variation coefficient is dimensionless. For non-negative random variables we have the relationship $\mathcal{E}\underline{x} = \int_0^\infty \mathcal{F}_{\underline{x}}(x) dx$ or $\mathcal{E}\underline{x} = \sum_{x=0}^\infty \mathcal{F}_{\underline{x}}(x)$.

Simple rules directly follow from the definitions, such as for a constant a , we must have $\mathcal{E}a = a$, $\text{var}(a\underline{x}) = a^2 \text{var } \underline{x}$, $\text{sd } a\underline{x} = a \text{sd } \underline{x}$, $\text{vc } a\underline{x} = \text{vc } \underline{x}$, $\text{var } a = 0$, $\mathcal{E}\underline{x}^2 \geq \mathcal{E}^2 \underline{x}$.

Two random variables \underline{x} and \underline{y} have a *simultaneous* (or joint) probability distribution $\Pr\{x_0 \leq \underline{x} \leq x_1, y_0 \leq \underline{y} \leq y_1\}$. For continuous random variables we can write $\Pr\{x_0 \leq \underline{x} \leq x_1, y_0 \leq \underline{y} \leq y_1\} = \int_{x=x_0}^{x_1} \int_{y=y_0}^{y_1} f_{\underline{x}, \underline{y}}(x, y) dx dy$. Two random variables are said to be stochastically *independent* if $\Pr\{x_0 \leq \underline{x} \leq x_1, y_0 \leq \underline{y} \leq y_1\} = \Pr\{x_0 \leq \underline{x} \leq x_1\} \Pr\{y_0 \leq \underline{y} \leq y_1\}$, i.e. $\int_{x=x_0}^{x_1} \int_{y=y_0}^{y_1} f_{\underline{x}, \underline{y}}(x, y) dx dy = \int_{x=x_0}^{x_1} f_{\underline{x}}(x) dx \int_{y=y_0}^{y_1} f_{\underline{y}}(y) dy$ for continuous random variables. For independent variables we have $\mathcal{E}\underline{x}\underline{y} = \mathcal{E}\underline{x} \mathcal{E}\underline{y}$.

The *covariance* between \underline{x} and \underline{y} is defined as $\text{cov}(\underline{x}, \underline{y}) = \mathcal{E}(x - \mathcal{E}x)(y - \mathcal{E}y) = \mathcal{E}\underline{x}\underline{y} - \mathcal{E}\underline{x} \mathcal{E}\underline{y}$, which is zero if \underline{x} and \underline{y} are independent. Notice that $\text{cov}(\underline{x}, \underline{x}) = \text{var } \underline{x}$ and $\dim(\text{cov}(\underline{x}, \underline{y})) = \dim(\underline{x}) \dim(\underline{y})$. The *correlation coefficient* between \underline{x} and \underline{y} is defined as

$$\text{cor}(\underline{x}, \underline{y}) = \frac{\text{cov}(\underline{x}, \underline{y})}{(\text{sd } \underline{x})(\text{sd } \underline{y})}.$$

The correlation coefficient can take values in the interval $(-1, 1)$, and equals 0 for independent variables. It is dimensionless. Notice that $\text{cor}(\underline{x}, \underline{x}) = \text{var } \underline{x} / (\text{sd } \underline{x} \text{sd } \underline{x}) = 1$.

Some implied rules for constant a and b : $\text{cov}(a\underline{x}, b\underline{y}) = ab \text{cov}(\underline{x}, \underline{y})$, $\text{cor}(a\underline{x}, b\underline{y}) = \text{cor}(\underline{x}, \underline{y})$. If \underline{x} and \underline{y} are independent, we have $\text{var}(\underline{x} + \underline{y}) = \text{var}(\underline{x}) + \text{var}(\underline{y})$, and $\text{var}(\underline{x} - \underline{y}) = \text{var}(\underline{x}) + \text{var}(\underline{y})$.

2.6.2 Examples of probability distributions

The *Poisson* probability distribution is defined by

$$\Pr\{\underline{x} = x\} = \frac{\lambda^x}{x!} \exp(-\lambda) \quad \text{for } x = 0, 1, \dots$$

where the parameter λ is non-negative. Properties are that $\mathcal{E}\underline{x} = \text{var } \underline{x} = \lambda$. Notice that λ , and therefore \underline{x} , must be dimensionless. The sum of independently Poisson distributed random variables is again Poisson distributed.

The *binomial* probability distribution is defined by

$$\Pr\{\underline{x} = x\} = \binom{n}{x} p^x (1-p)^{n-x} \quad \text{for } x = 0, 1, \dots, n$$

where the parameter p can take values in the interval $(0, 1)$, and the parameter n can take the values $\{1, 2, \dots\}$. Properties are $\mathcal{E}\underline{x} = np$, and $\text{var } \underline{x} = np(1-p)$. The parameters p and n , and \underline{x} are dimensionless. The binomial distribution converges to the Poisson distribution for $np \ll n$. For $n = 1$, the binomial distribution is also called the *Bernoulli distribution*. If x_i follows a Bernoulli distribution with parameter p , $\sum_{i=1}^n x_i$ follows a binomial distribution with parameters p and n .

The *multinomial* probability distribution is an extension of the binomial one, and defined as

$$\Pr\{\underline{x}_1 = x_1, \dots, \underline{x}_s = x_s\} = \frac{n!}{x_1! \dots x_s!} p_1^{x_1} \dots p_s^{x_s} \quad \text{for } x_i = 0, 1, \dots, n \text{ and } \sum_i x_i = n$$

with $\sum_{i=1}^s p_i = 1$. The parameter s takes values $\{2, 3, \dots\}$, and x_i takes values $\{0, 1, \dots, n\}$, and p_i takes values in $(0, 1)$, for all $i = 1, \dots, s$. Notice that the multinomial distribution reduces to the binomial one for $s = 2$; since $\underline{x}_2 = n - \underline{x}_1$, it is no longer of much interest, and we call \underline{x}_1 just \underline{x} , and suppress reference to \underline{x}_2 .

The *geometric* probability distribution is defined as

$$\Pr\{\underline{x} = x\} = p(1-p)^x \quad \text{for } x = 0, 1, \dots$$

where parameter $0 < p < 1$ has the interpretation of the probability of success when \underline{x} measures the number of unsuccessful trials before the first successful one if the trials are independent. We have $\mathcal{E}\underline{x} = (1-p)/p$ and $\text{var } \underline{x} = (1-p)/p^2$. If \underline{x} measures the number of unsuccessful trials before the r -th successful one (so the r -th success occurs at trial $r + x$), \underline{x} is *negative binomially* distribution

$$\Pr\{\underline{x} = x\} = \binom{x+r-1}{x} p^r (1-p)^x \quad \text{for } x = 0, 1, \dots$$

If x_i is geometrically distributed with parameter p , $\sum_{i=1}^r x_i$ is negative binomially distributed with parameters p and r . We have $\mathcal{E}\underline{x} = r(1-p)/p$ and $\text{var } \underline{x} = r(1-p)/p^2$.

2.6.3 Examples of probability density functions

The *homogeneous* distribution on the interval (a, b) is given by

$$f_{\underline{x}}(x) = (x > a)(x < b) \quad \text{for } -\infty < x < \infty$$

The mean in $\mathcal{E}\underline{x} = (a+b)/2$ and variance $\text{var } \underline{x} = (b-a)^2/12$. We have $\dim(a) = \dim(b) = \dim(x)$. If \underline{x} is homogeneously distributed on the interval $(0, 1)$, $-\log \underline{x}$ is exponentially distributed with parameter 1. The importance of this fact is the readily availability of random generators for homogeneously distributed random variables, which are used in Monte Carlo methods.

The *beta* distribution is defined by

$$f_{\underline{x}}(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)} \quad \text{for } 0 \leq x \leq 1$$

which has mean $\mathcal{E}\underline{x} = \frac{\alpha}{\alpha+\beta}$ and variance $\text{var } \underline{x} = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$. Both α , β and x must be dimensionless.

A random variable is said to be *exponentially* distributed if

$$f_{\underline{x}}(x) = \lambda \exp(-\lambda x) \quad \text{for } x \geq 0$$

where λ and \underline{x} take values in $(0, \infty)$, while $\dim(\lambda) = 1/\dim(x)$. We have $\mathcal{E}\underline{x} = \lambda^{-1}$, $\text{var } \underline{x} = \lambda^{-2}$, $F_{\underline{x}}(x) = 1 - \exp(-\lambda x)$, $\mathcal{F}_{\underline{x}}(x) = \exp(-\lambda x)$. The relationship between the exponential and the Poisson distribution is explained in the section on point processes.

The sum of n independent exponentially distributed random variables with parameter λ is *gamma* distributed:

$$f_{\underline{x}}(x) = \frac{\lambda}{\Gamma(n)} (\lambda x)^{n-1} \exp(-\lambda x) \quad \text{for } x \geq 0$$

which has mean $\mathcal{E}\underline{x} = n/\lambda$ and variance $\text{var } \underline{x} = n/\lambda^2$. We have $\dim(\lambda) = 1/\dim(x)$, and n is dimensionless.

The uni- and p -variate *normal* p.d.f. are defined as

$$\begin{aligned} f_{\underline{x}}(x) &= (2\pi\sigma^2)^{-1/2} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \quad \text{for } -\infty < x < \infty \\ f_{\underline{x}}(x) &= (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)\right) \quad \text{for } -\infty < x_i < \infty \end{aligned}$$

where μ takes values in $(-\infty, \infty)$, and σ and Σ in $(0, \infty)$. Properties are that $\mathcal{E}\underline{x} = \mu$ and $\text{var } \underline{x} = \sigma^2$, which shows that $\dim(\mu) = \dim(\sigma) = \dim(x)$. The significance of the normal distribution is by the central limit theorem, which states that the sum of independent identically distributed random variable is asymptotically normally distributed, for large enough number of random variables (irrespective of their own distribution).

The χ^2 p.d.f. is defined as

$$f_{\underline{x}}(x) = (2^{\nu/2} \Gamma(\nu/2))^{-1} x^{\nu/2-1} \exp(-x/2) \quad \text{for } x \geq 0$$

The dimensionless parameter ν is usually called the degree of freedom, and can assume a value in the interval $(0, \infty)$. Notice that $\Gamma(\nu + 1) = \nu!$, for $\nu = 0, 1, \dots$. We have that $\mathcal{E}\underline{x} = \nu$. The sum of ν squared independent normally distributed random variables with mean 0 and variance 1 is χ^2 -distributed with ν degrees of freedom. Moreover, if $\underline{x} \equiv \{\underline{x}_1, \underline{x}_2, \dots, \underline{x}_p\} \sim N(\underline{\mu}, \Sigma)$, the quadratic form $(\underline{x} - \underline{\mu})' \Sigma^{-1} (\underline{x} - \underline{\mu})$ is χ^2 distributed with p degrees of freedom.

2.6.4 Conditional and marginal probabilities

Suppose the \underline{x} can assume values in Ω . The probability on the result $\underline{x} = x$ for some $x \in \Omega_0$, given that \underline{x} assumes values in Ω_0 , while Ω_0 is a subset of Ω , equals

$$\Pr\{\underline{x} = x | \underline{x} \in \Omega_0\} = \Pr\{\underline{x} = x\} / \Pr\{\underline{x} \in \Omega_0\}$$

This is called a *conditional probability*. If \underline{x} follows a Poisson distribution, for example, the probability of \underline{x} , given that it is larger than 0 is given by $\Pr\{\underline{x} = x | \underline{x} > 0\} = \Pr\{\underline{x} = x\} / (1 - \Pr\{\underline{x} = 0\}) = \frac{\lambda^x}{x!} \exp(-\lambda) / (1 - \exp(-\lambda))$. This is an example of a *truncated* probability distribution.

Likewise, if \underline{x} and \underline{y} follow some simultaneous probability distribution, the probability that \underline{x} attains some value, given that \underline{y} has some specified value equals

$$\Pr\{\underline{x} = x | \underline{y} = y\} = \Pr\{\underline{x} = x, \underline{y} = y\} / \Pr\{\underline{y} = y\}$$

where $\Pr\{\underline{y} = y\} = \sum_x \Pr\{\underline{x} = x, \underline{y} = y\}$ is the *marginal distribution* of \underline{y} . For continuous random variables \underline{x} and \underline{y} we have the *conditional p.d.f.* $f_{\underline{x}}(x | y) = f_{\underline{x}, \underline{y}}(x, y) / f_{\underline{y}}(y)$, where $f_{\underline{y}}(y) = \int_x f_{\underline{x}, \underline{y}}(x, y) dx$ is the *marginal p.d.f.* of \underline{y} . The conditional p.d.f. of \underline{x} has dimension $\dim(x)^{-1}$.

If \underline{x} follows some probability distribution with parameter θ , while this value θ represents a random trial from some p.d.f. we have the *mixture* $\Pr\{\underline{x} = x\} = \int_{\theta} \Pr\{\underline{x} = x | \theta\} f_{\theta}(\theta) d\theta$.

2.6.5 Calculations with random variables

The probability distribution of a *sum* \underline{z} of two random variables \underline{x} and \underline{y} is given by $\Pr\{\underline{z} = z\} = \sum_v \Pr\{\underline{x} = v, \underline{y} = z - v\}$. If \underline{x} and \underline{y} are independent, this reduces to $\Pr\{\underline{z} = z\} = \sum_v \Pr\{\underline{x} = v\} \Pr\{\underline{y} = z - v\}$. Likewise for independent continuous random variables we have $f_{\underline{z}}(z) = \int_v f_{\underline{x}}(v) f_{\underline{y}}(z - v) dv = \int_v f_{\underline{x}}(z - v) f_{\underline{y}}(v) dv$. Such an integral is known as a *convolution integral*.

The probability distribution of a *product* \underline{z} of two random variables \underline{x} and \underline{y} is given by $\Pr\{\underline{z} = z\} = \sum_v \Pr\{\underline{x} = v, \underline{y} = z/v\}$; for independent continuous random variables we have $f_{\underline{z}}(z) = \int_v f_{\underline{x}}(v) f_{\underline{y}}(z/v) dv$.

The survivor function of the *maximum* \underline{y} of n independent identically distributed random variables $\{\underline{x}_i\}_{i=1}^n$ is $S_{\underline{y}}(y) = S_{\underline{x}}(y)^n$. The distribution function of the *minimum* \underline{y} of n independent identically distributed random variables $\{\underline{x}_i\}_{i=1}^n$ is $F_{\underline{y}}(y) = F_{\underline{x}}(y)^n$.

Suppose that $\underline{y} = g(\underline{x})$ is some monotonous *transformation* g of \underline{x} , with inverse G , so $x = G(g(x))$. If \underline{x} has p.d.f. $f_{\underline{x}}(x)$, the distribution function of \underline{y} is $\Pr\{\underline{y} > y\} = \Pr\{\underline{x} >$

$G(y)$ and the p.d.f. of \underline{y} is $f_{\underline{y}}(y) = f_{\underline{x}}(G(y))\frac{d}{dx}g(x)$. If \underline{x} is homogeneously distributed on $[0, 1]$, so $f_{\underline{x}}(x) = (x \geq 0)(x \leq 0)$, and if $g(x) = -\log x$, then $G(y) = \exp(-y)$, and $f_{\underline{y}}(y) = (y > 0)\exp(-y)y^{-1}$; so \underline{y} is exponentially distributed.

2.7 Numerical methods

The art of numerical methods is to find numerical values for quantities that cannot be obtained analytically. The problem is always to specify the desired accuracy, in combination with robustness and “reasonable” computational effort. Many numerical procedures have been developed, we only illustrate a very few simple ones to illustrate the concepts. The rapid increase of the computer performance makes that computational effort is continuously changing in appreciation, which makes that numerical methods are still in full development. For simple applications it has become close to meaningless, but for many problems computational effort is still an issue to take into consideration. This is also caused by the fact that problems are nowadays solved numerically, that were avoided in earlier days for obvious reasons.

2.7.1 Numerical integration

The trapezoidal approximation for a definite integral is

$$\int_a^b f(x) dx \simeq \frac{\Delta x}{2}(y_0 + 2y_1 + \cdots + 2y_{n-1} + y_n)$$

where $\Delta x = (b - a)/n$ for some appropriate choice of n , and $y_i = f(a + i\Delta x)$. Tangent line and tangent parabola approximations usually work better for smooth functions, i.e. the balance between accuracy and computational effort.

For a differential equation $\frac{d}{dt}x = f(t, x)$, the second order Runge-Kutta method reads

$$x_{n+1} = x_n + (k_1 + k_2)/2 + O(h^3) \quad \text{for } k_1 = hf(t_n, x_n) \quad \text{and } k_2 = hf(t_n + h, x_n + k_1)$$

for some chosen step size h , with $t_{n+1} = t_n + h$. Starting from $(0, x_0)$, the pairs (t_i, x_i) are generated till some chosen endpoint t_N , while x may be vector-valued. If f does not depend on x , the solution of the differential equation reduces to plain integration, and the second-order Runge-Kutta method reduces to the trapezoidal approximation. Higher order Runge-Kutta schemes are devised; the fifth-order allows a variable step size, where the steps are small if the function f changes rapidly. Differential equations can behave in *stiff* way, which means that the values change slowly for most of the time, but sometimes they do change rapidly. The integration of such equations require special techniques, because the rapid changes force a small choice of step size h , which is most of the time not necessary.

See also splines at {53}.

2.7.2 Numerical differentiation

Numerical differentiation follows the definition of differentiation, omitting the limit. This makes that the forward and backward incremental step in the independent variable no




Figure 2.3: The Newton Raphson procedure for finding a root of a function of a single variable. The root of the tangent line at a point is used as the new point. The approximate domain of attraction is given in green. Any choice for a starting point x_0 outside this domain will not result in an approximation for the root.

longer has the same result. We take the mean with

$$\frac{d}{dx}f(x_1) \simeq \frac{f(x_1 + \Delta x) - f(x_1 - \Delta x)}{2\Delta x}$$

for some appropriate choice of Δx . Numerical differentiation easily gives substantial deviations from the real thing; whenever possible it should be avoided. This even more applies to the numerical approximation of the second derivative

$$\frac{d^2}{dx^2}f(x_1) \simeq \frac{f(x_1 + \Delta x) - 2f(x_1) + f(x_1 - \Delta x)}{(\Delta x)^2}.$$

See also splines at {53}.

2.7.3 Root finding

Newton Raphson's iteration scheme for finding the roots of equation $f(x) = 0$ is

$$x_{n+1} = x_n - \Delta x_n \quad \text{with} \quad \Delta x_n = \left(\frac{d}{dx^T} f(x_n) \right)^{-1} f(x_n)$$

where both f and x can be vector valued. Starting from an appropriately chosen x_0 , the convergence is usually very fast, but the basin of attraction is rather small (so x_0 has to be close to the real root). The iteration is stopped when the norm $\|f(x)\|$ is smaller than a specified value, or the number of iterations exceeds a maximum. The basin can usually be extended by restricting the step size $\|\Delta x\|$ to some upper limit by multiplication of Δx with a scalar. Be aware of the problem of multiple roots. The iteration only finds one, which depends on the starting value x_0 , see Figure 2.3.

Newton's method is frequently used for finding extremes, by applying it to the derivatives of the function. This only works if the function satisfies smoothness conditions. Be aware of the problem that Newton's method only finds a critical point, which is not always the one you want to have. The result requires testing to tell critical points apart from extremes (and minima from maxima). Several starting values should be used to find other critical points, but one cannot be sure that all critical points are detected.

See also splines at {53}.

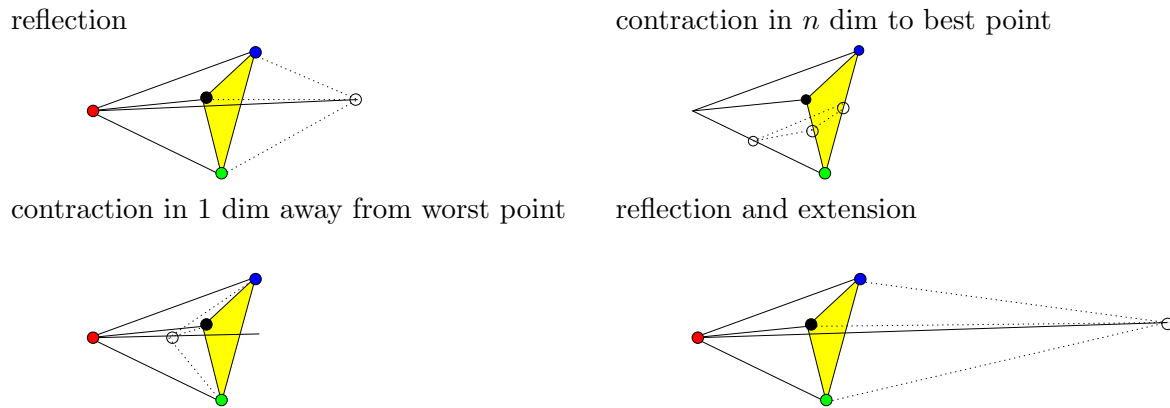


Figure 2.4: The simplex method in 3 variables, so the simplex consists of 4 points. The worst point (red) is replaced by a better point (open), according to one of 3 possible strategies, or the simplex is contracted to the best point (green). The second best point is blue. The vertices of the new simplex are connected with dotted lines.

2.7.4 Extreme finding

The downhill simplex method is due to Nelder and Mead [?]. It is slow, but also very robust; it requires no derivatives. If the minimization is in n variables, simplex method starts in $n + 1$ points (or vertices), and the worst point is replaced by a better point according to the scheme indicated in figure 2.4. A procedure is stopped when the vertices are closer together than some threshold value, or the function values at the vertices differ less than some threshold value, or the number of function evaluations exceeds some threshold value (unsuccessful termination).

Chapter 3

Models for processes

Life is essentially a process in time, no wonder that the analysis of processes is basic in biology. Two aspects of processes are usually studied: *transient states*, where the process is followed while it evolves in time, and *steady states*, where the long-term behaviour is studied, that is independent of how the process originally started.

3.1 Types of processes

Processes can be modeled such that time can take particular values only (*discrete time processes*, where the distance between the time points, the time steps, is usually constant) or any value (*continuous time processes*)

The specification of discrete time models is usually done in the form of *maps*, where the value of quantities at time t is written as some function of the values at earlier points in time.

Most models for processes are *deterministic*, which means that they do not include random variables, contrary to *stochastic* models. One of the main reasons is that realistic implementations of random variables very easily lead to extremely complex models. Most stochastic models (such as the class of stochastic differential equations) do not satisfy e.g. laws for energy and mass conservation in the strict sense. Their statistical evaluation is very complex, so that one usually has to rely on computer simulation studies. We here only discuss Markovian models, which still have some basic simplicity. Deterministic models for processes obviously suffer from the problem that they are hardly realistic for since stochastic behaviour is basic to most biological processes.

3.1.1 Stochastic processes

Correlation functions

If the pairs $\{\underline{x}(t_i), \underline{y}(t_i)\}_i$ form a stochastic process with constant time step $\Delta t = t_{i+1} - t_i$, the function $\rho_{\underline{xy}}(h\Delta t) = \text{cor}(\underline{x}(t_i), \underline{y}(t_{i+h}))$ is known as the *cross-correlation* function, and the $\rho_{\underline{xx}}(h\Delta t) = \text{cor}(\underline{x}(t_i), \underline{x}(t_{i+h}))$ as the *auto-correlation* function. They quantify the memory of the process in a particular way. Similar constructs are defined for continuous

time, and for more than two variables. Correlation functions are matrix-valued, where the independent variable represents the time shift between the members of correlated pairs.

Markov chains

A *Markov chain* is a model of the type

$$p(t_{n+1}) = Pp(t_n) \quad (3.1)$$

where $p(t_n)$ is the vector with probabilities that the system is in its various states, and P the square matrix with fixed transition probabilities; the typical element p_{ij} gives the probability that the system will be in state i in the next time step, given that it is in state j . By definition we must have that $P\mathbf{1} = \mathbf{1}$ and $p(t_i)^T\mathbf{1} = 1$, where $\mathbf{1}$ is a vector with ones.

A consequence of the construct is that $p(t_{n+r}) = P^r p(t_n)$, and that the steady state probability distribution satisfies $Pp(t_\infty) = p(t_\infty)$, in other words it is an eigen vector of P that is associated with eigen value 1.

The simplism of the Markov chain derives from the fact that it has no memory; it does not matter what path the system took to arrive in a certain state. Some memory can be built into the Markov chains by including the path in the definition of the state. So if the system has originally n states, we delineate n^2 states to include the previous state, and n^3 states to include the last two states.

A *branching process* is a Markov chain of a special structure, where the state of the system represents the total number of objects and at each time step each object is replaced by a number that represents a random trial from some probability distribution. This is typical for particular types of organisms, but interactions among organism are difficult to implement in branching processes.

Markov processes

Markov processes are Markov chains in continuous time. The transition from discrete to continuous time can be smooth for a particular type of Markov chains, where the states represent numbers in a population of objects. Transition probabilities in an infinitesimally small time increment are only positive for neighbouring states, and proportional to the time increment. So we arrive at the specification for $dt \downarrow 0$

$$\Pr\{\underline{x}(t+dt) = m | \underline{x}(t) = n\} = \begin{cases} \lambda dt & \text{if } m = n - 1 \\ 1 - (\lambda + \mu) dt - o(dt) & \text{if } m = n \\ \mu dt & \text{if } m = n + 1 \\ o(dt) & \text{if } m < n - 1 \quad \text{or} \quad m > n + 1 \end{cases} \quad (3.2)$$

where $\lim_{dt \downarrow 0} o(dt)/dt = 0$ by definition. This specification can be rewritten in the differential equation for $p_n(t) \equiv \Pr\{\underline{x}(t) = n\}$

$$p'_n(t) = \mu p_{n-1}(t) + \lambda p_{n+1}(t) - (\lambda + \mu) p_n(t)$$

Together with a specification of the initial conditions, such a process is called a *birth and death process*. For $\lambda = 0$ and $\mu = x\beta$, for a constant β , the process is called the

Yule process. The birth rate is here proportional to the population size, which is typical for organisms. The inclusion of interactions among organisms (such as competition for resources), however, is difficult to implement in Markov processes.

Point processes

A *point process* is a special type of stochastic process, where the time points of the occurrence of events are stochastic. Such processes have two aspects: the *counting process*, which follows the number of events in a fixed time interval, and the *interval process*, which deals with the time intervals between subsequent events. If the probability of two events in a single time interval is negligibly small, we call the process *orderly*.

If subsequent time intervals represent independently identically distributed random variables, we call the process a *renewal process*. The term renewal process originates from a special type of point process, where the events represent replacement of an object, such as a light bulb for instance. If these intervals are independently exponentially distributed, we call the process a *Poisson process*, because the number of events in a time interval is then Poisson distributed.

The *intensity* of a point process is defined as the expected number of events per time increment, considered as a function of time. If the intensity is constant, we call the process *ergodic*. The intensity equals the *hazard rate*. The hazard rate is defined as

$$h_{\underline{t}}(t) = \frac{f_{\underline{t}}(t)}{S_{\underline{t}}(t)},$$

where the continuous random variable \underline{t} frequently has the interpretation of the life span of an object. It is also called mortality rate; the product $h_{\underline{t}}(t) dt$ specifies the probability that the object will die in the infinitesimally small time interval $(t, t + dt)$, given that it is still alive at t . So it is a conditional p.d.f.. Notice that $h_{\underline{t}}(t) = -\frac{d}{dt} \ln(S_{\underline{t}}(t))$, so $S_{\underline{t}}(t) = \exp\{-\int_0^t h_{\underline{t}}(s) ds\}$.

If the hazard rate is constant, the process is a Poisson process, so the number of death in a fixed interval is Poisson distributed, and the intervals between subsequent deaths are exponentially distributed. A central theorem states that the sum of an increasing number of mutually independent ergodic point processes converges to a Poisson process.

3.2 Systems

The idea behind the concept of a system is simple in principle. A system is based on the idea of *state variables*, which are supposed to specify completely the state of the system at a given moment. Completeness is essential. The next step is to specify how the state variables change with time as a function of a number of *inputs* and each other. The specification usually takes the form of a set of *ordinary differential equations* (ode' s)

$$\frac{d}{dt}x = f(x|\theta) \quad \text{for } x = (x_1, \dots, x_n)^T \quad (3.3)$$

which have parameters θ , i.e. constants that are assumed to have some fixed value in the simplest case. Usually this specification also includes a number of *outputs*. The set of differential equations fully specifies the dynamics of the system in combination with the specification of the system at the start, $x(0)$, or at some moment in time, $x(t_1)$. Finding the states of the system as a function of time is called an *initial value* problem, of a *boundary value* problem, respectively.

Parameters are typically constant, but sometimes the values change with time. This can be described by a function of time, which again has parameters that are now considered to be constant. For instance, parameters that have the interpretation of physiological rates depend on temperature; therefore, they remain constant as long as the temperature does not change. If the temperature does change, then the parameters do as well. Heat, however, is generated as a side product of metabolism. In ectotherms, i.e. animals that do not heat their body to a constant high temperature, heat production is low, because of their usually low body temperature. The body temperature usually follows that of the environment, and can thus be treated as a function of time. The situation is more complex in developing birds, which make the transition to the endothermic state some days after hatching. The hatchling's temperature is high, because of brooding; therefore, metabolism and heat production are also high. In addition, the young bird starts to invest extra energy in heating. Here, the state variables of the system interfere with the environment, but not via input; this means that the body temperature must be considered as an additional state variable.

Choosing the state variables is the most crucial step in defining a system. It is usually a lot easier to compare and test alternative formulations for the change of state variables, than different choices of state variables. Models with different sets of state variables are hardly comparable, both conceptually and in tests against data. Statistics basically deals with parameter values, and is of little use when comparing the goodness of fit of models that differ in structure.

3.2.1 Constraints on dynamics

Mass, energy and other conserved quantities pose constraints on the possible behaviour of dynamic systems. The explicit use of these constraints forms a powerful tool in the specification of the dynamics of the system. If the states are appropriately specified, these k constraints on x_1, \dots, x_n can be written in the form $g_i(x) = 0$, for $i = 1, \dots, k$. If all x_i represent masses, the functions g_i take the form $\sum_{ij} w_{ij} x_j = 0$, where w_{ij} are fixed weight coefficients. The system can be reduced to $n - k$ variables, and we have to find the remaining k variables from the k constraints. If the system is not reduced, the Jacobian matrix of the system will have k eigen values equal to zero at any point in time. This might hamper the application of software for the analysis of asymptotic properties of the system.

3.2.2 Feedback

Systems can behave in ways that cause an amplification of that behaviour, a phenomenon called *positive feedback*, or a reduction of that behaviour, called *negative feedback*. A growing population of organisms is likely to grow faster because more organisms partake in reproduction (positive feedback), but they exhaust their resources sooner (negative feedback). The notion “feedback” originates from engineering, where systems are constructed. In biology, where systems are usually given, and many components of the system work in opposite directions, the notion is less operative.

3.2.3 Asymptotic properties

Some model systems have infinite *memories*. Think, e.g. of a deterministic model for cell growth and division in a homogeneous environment where each cell divides into two identical daughter cells when it reaches a certain size. Since all cells experience the same environment, all daughter cells of some cell in the inoculum (i.e. the founder population) will divide at the same moment, because they all experience the same environment. Such a system remains dependent on properties of the inoculum (e.g. the size distribution). Such (unrealistic) memory can be removed in different ways (such as small random differences in size between daughters, or individual-specific size thresholds for division), but this makes the model more complex. We here discuss asymptotic properties of models that have finite memories, so all information about the (remote) past disappears, and focus on properties of a set of ODEs of the type $\frac{d}{dt}x = f(x)$, for some vector-valued x , and how it depends on the parameters of the system..

The first important observation is that, if we follow $x(t)$ through time by evaluating $x(t) = \int_0^t f(x(s)) ds$, starting from some $x(0)$, several things can happen. First the system can become *degenerated*, i.e. one or more components of x can become zero, or $\pm\infty$ (the latter possibility is usually excluded by conservation laws). If not, it can evolve to some constant value x^* , or it can start to change cyclically. This behaviour might depend on the choice for $x(0)$. The system is “attracted” by a *point attractor* or a *cyclic attractor*. When the number of loosely coupled variables is sufficiently large in a system, the system is likely to have very complex asymptotic behaviour for some combination of parameter values, including the occurrence of multiple attractors, possibly of the chaotic type, called *strange attractors*. A deterministic system is said to behave chaotically if it varies in time without repeating itself periodically. Since this is difficult to check (because periods can be extremely large), another property is used in practice: an infinitesimally small change in initial values of the system eventually results in a substantial difference in behaviour. This property is not unique for chaos, however. The behaviour of the system is usually compared with that for slightly different values of particular parameters of the system, to see if the pattern matches the known ‘routes to chaos’. It remains difficult to be sure that a system behaves chaotically, but it is not rare to find chaos in the more complex models. This is almost independent of the specific model.

Many asymptotic properties of a set of ODEs $x' = f(x)$ can be deduced from the eigenvalues of the Hessian $\frac{\partial^2}{\partial x \partial x^T} f(x^*)$ evaluated at an equilibrium x^* of the system $f(x^*) = \mathbf{0}$.

Table 3.1: List of basic local bifurcations for ODEs: $dx/dt = f(x, \alpha)$, and maps: $y_{n+1} = f(y_n, \alpha)$ with normal forms. The bifurcation point is $\alpha = 0$. λ is the eigenvalue of the Hessian evaluated at the equilibrium ($\frac{d}{dt}x = 0$ and $y_{n+1} = y_n$) and μ is the Floquet multiplier evaluated at the limit cycle. The bifurcation type depends on the real (Re) parts of these characteristic exponents. A stable positive attractor originates at a supercritical transcritical bifurcation (superscript +) and an unstable positive equilibrium or limit cycle at a subcritical transcritical bifurcation (superscript -). Superscript \pm refers to supercritical and subcritical.

symbol	bifurcation	normal form	characteristic exponents
T_e	Tangent, of equilibrium	$\frac{d}{dt}x = \alpha - x^2$	Re $\lambda = 0$
T_c	Tangent, of limit cycle	$y_{n+1} = y_n + \alpha - y_n^2$	Re $\mu = 1$
TC_e^\pm	Transcritical, of equilibrium	$\frac{d}{dt}x = \alpha x \pm x^2$	Re $\lambda = 0$
TC_c^\pm	Transcritical, of limit cycle	$y_{n+1} = (1 + \alpha)y_n \pm y_n^2$	Re $\mu = 1$
H^\pm	Hopf	$\frac{d}{dt}x = -y + x(\alpha \pm (x^2 + y^2))$ $\frac{d}{dt}y = x + y(\alpha \pm (x^2 + y^2))$	Re $\lambda_{1,2} = 0$
F^\pm	Flip	$y_{n+1} = -(1 + \alpha)y_n \pm y_n^3$	Re $\mu = -1$

If such an equilibrium exists for non-zero values of the variables and the eigen values are all real, the equilibrium represents a point attractor; if some are complex, the equilibrium is unstable, and has a limit cycle associated with it. The imaginary part gives information about the orbit. If a system happens to be at an unstable equilibrium it will remain there for ever, but if it is slightly disturbed (i.e. placed out of equilibrium) it will never return to an unstable equilibrium. If a system is slightly disturbed while it was in a limit cycle, it may return to that limit cycle, and we then speak of a stable limit cycle, or it might follow a different limit cycle, in which case we speak of *neutral stability*. The study of the reactions of a system to small perturbations is called *local stability* analysis, which does not give information about reactions to large perturbations, called *global stability* analysis. A system can have more than one attractor.

In practice we first try to find an approximation for a point attractor by following the system through time (e.g. using a Runge Kutta method), and then we try to find the roots of the system $f(x^*) = 0$ (e.g. using a Newton Raphson method). Once we have found a root, we can change one or more parameters a little, and find a new root with the Newton Raphson method, using the previous roots as a starting value. This technique is called *continuation*. Notice that we can find unstable roots this way, to which the system cannot evolve.

Bifurcation analysis deals with qualitative changes in the asymptotic behaviour of the system, when a parameter is varied in value; this selection from the list of parameters of the system is called bifurcation parameter. Its interest rests on the interpretation of the system, so no general rules can be given for its selection. Table 3.1 gives the possible bifurcation types. The bifurcation type depends on the value of the eigenvalue of the Hessian matrix evaluated at the equilibrium and the Floquet multiplier, which is an eigenvalue of the Poincaré next-return map. If all complex values of the Floquet multipliers are within the unit circle, the dynamic system's orbit converges to a limit cycle.

The analysis of bifurcation behaviour must be done numerically, using specialized software: LOCBIF [?] and AUTO [?] can calculate bifurcation diagrams using continuation methods. The theory is documented in [?]. The analyses cannot be done on a routine basis, however, and the user must have a fairly good idea of what to expect and what to look for. Although the software is rapidly improving in quality, at present it is still deficient in computing certain types of global bifurcations, for instance, and one has to rely on ‘in-house’ software to fill in the gaps, see [?].

Results of bifurcation analyses are frequently reported in the form of bifurcation diagrams. These diagrams connect points where system’s asymptotic behaviour changes in a similar way when the bifurcation parameters are varied. So, the system has similar asymptotic behaviour for values of bifurcation parameters within one region. The construction of such diagrams is only feasible if there are just one or two of such parameters. Diagrams with these parameters are called operating diagrams.

Chapter 4

Model-based statistics

4.1 Scope of statistics

Statistics deals basically with the estimation of model parameters from data and with the testing of hypotheses about their values; do comparable parameter values in two samples differ significantly, or does a particular parameter value differ significantly from some given value? These ideas may be used for optimizing experimental design and sampling strategies.

Statistics cannot deal with problems like: is this model “correct”, or does this model fit significantly better than that model? Statistics treats the model as given. The goodness of fit of a model can be quantified. Only stochastic models can be tested against data; deterministic models are usually extended to stochastic ones, by introducing a ‘measurement error’, and treating it as a regression model. This is convenient, but rarely realistic. Models might fit data well for the wrong reasons; models are idealizations of reality, so we can only expect some deviations from model predictions. To what extent deviations from model predictions are problematic depends on the purpose, so on the context. Since we deal with probabilities, we cannot be sure of anything while using statistics; a correct model might fit data poorly.

A model gets its value from the mechanistically inspired assumptions from which it is derived. Without such assumptions (so if the model itself is the assumption), the model is close to useless, including all statistical inference based on it. This is the reason why *one should never transform data*, wherever statistical text books might tell about this; transformations destroy the relationship between the model and its assumptions, and so the usefulness of the model.

Many statistical methods are based on linear models (e.g. ANOVA, multiple correlation, principal component analysis, factor analysis, auto-regression). Since such models rarely apply in biology, these methods are not discussed here.

4.2 Measurements: scales and units

A *measurement* assigns a numerical value to a quantity (object or process). Depending on the nature of this quantity, the measurement can be in one of the following *scales*:

a	annum of time (1 a \simeq 365.25 d = 31.56 Ms)	A	ampere of electric current
At	ampere-turn of magnetomotive force	C	coulomb of electrical charge (1 C = 1 A S)
$^{\circ}\text{C}$	degree Celsius ($0^{\circ}\text{C} = 273.15\text{ K}$)	cd	candela of luminous intensity
d	day of time (1 d = 24 h = 86.4 ks)	F	farad of electric capacitance (1 F = 1 C V $^{-1}$)
g	gram of mass	h	hour of time (1 h = 3600 s)
H	henry of inductance (1 H = 1 Wb A $^{-1}$)	Hz	hertz of frequency (1 Hz = 1 s $^{-1}$)
J	joule of energy (1 J = 1 N m)	K	kelvin of temperature
l	liter of volume (1 l = 1 dm 3)	lm	lumen of luminous flux (1 lm = 1 cd sr $^{-1}$)
lx	lux of illumination (1 lx = 1 lm m $^{-2}$)	m	meter of length
mol	mole of compound (1 mol = 6.02 10 23 molecules)	nt	nit of luminance
N	newton of power (1 N = 1 kg m s $^{-2}$)	Ω	ohm of resistance (1 Ω = 1 V A $^{-1}$)
Pa	pascal of pressure (1 Pa = 1 N M $^{-2}$)	rad	radian of plane angle
s	second of time	S	siemens of electrical conduction (1 S = 1 Ω^{-1})
sr	steradian of solid angle	T	tesla of magnetic flux density (1 T = 1 Wb m $^{-2}$)
V	volt of potential difference (1 V = 1 W A $^{-1}$)	W	watt of power (1 W = 1 J s $^{-1}$)
Wb	weber of magnetic flux (1 Wb = 1 V s)		

Table 4.1: Symbols for important single units of measurement in the SI system.

Nominal scale The numbers only represent a name for a category of objects, e.g. sexes of organisms can be numbered as 0, 1, \dots (yes, some organisms have 4 sexes). No operators are defined for this scale, and the scale is invariant under bijection.

Ordinal scale The numbers have an order, such as marks of an exam, or wind speed in Beaufort. Addition and multiplication are not defined for this scale. This scale is invariant under monotone transformation.

Interval scale The difference between numbers have a meaning, such as temperature in degrees Celsius or Fahrenheit. This scale is invariant under linear transformation ($y = ax + b$).

Ratio scale The difference and ratio's between numbers have a meaning, which implies that the value zero is not arbitrary, such as temperature in Kelvin, or speed. This scale is invariant under proportional transformation ($y = ax$).

Absolute scale Like the ratio scale, so the scale has a natural origin, but it also has a natural unit. The classic example is counts.

Numbers on the interval and ratio scales have *units*. These units standardized in the International System, and symbols are associated with these standardized units, see Table 4.1.

To avoid the notation of small or large numbers, submultiple prefixes are used, such as in ms (milli-second) or km (kilo-meter); see Table 4.2.

4.3 Precision and accuracy

The precision with which a variable should be measured completely depends on the aim of the measurement. Before measuring, we therefore ought to have an idea of what to

10^{-18}	a	atto-	10^{-15}	f	femto-	10^{-12}	p	pico-	10^{-9}	n	nano-
10^{-6}	μ	micro-	10^{-3}	m	milli-	10^{-2}	c	centi-	10^{-1}	d	deci-
10^1	da	deka-	10^2	h	hecto-	10^3	k	kilo-	10^6	M	mega-
10^9	G	giga-	10^{12}	T	tera-	10^{15}	P	peta-	10^{18}	E	exa-

Table 4.2: Standardized prefixes that can be used in combination with SI units.

do with the result. Too high an precision, or too low can be a waste of effort. The circumstance that the one planning the measurement is often somebody else than the one actually performing the measurement should stimulate an explicit statement of the desired precision. The presentation of the measurement should reflect its precision in the form of the correct number of so-called *significant figures*. This number is just one larger than the number of which one is certain. The number 13.40, or 0.001340, has 4 significant figures. The measurer is certain of the first three figures, but the last one, 0, is estimated. Non-significant zero's are always suppressed. Therefore, the number of significant figures in 134000 is not obvious. It can range from 3 to 6. In order to make the number of significant figures explicit and, sometimes, to shorten notation, the notion of *floating points* has been introduced. In a floating point notation, one places the point just behind the first figure and multiply with the appropriate integer power of 10. So 1.34 104 has 3 significant figures, while 1.3400 104 has got 5. Computer related manuscripts frequently use 'E' for 'times ten to the power', like 1.34E4 or 1.3400E4.

By *rounding off*, one reduces the number of significant figures: One replaces the number by the nearest number with one (or more) significant figures less. The number 14.0979 is rounded to 14.098, 14.10, 14.1, 14, and 1E1 depending on the desired number of significant figures. In calculations, one should only round numbers in the end result only, and not in intermediate results. Otherwise, errors can build up rapidly. In machines, figures become lost by a process called *truncation*, where they are simply deleted. Usually, the number of figures is sufficiently large to be of no practical problem. Sometimes, however, it is not always easy to tell when, errors induced by this truncation can built up to an intolerable extend. In such case one should change algorithm or machine.

Measurements can usually not be repeated exactly. The deviation from the unknown 'true' value, is termed error, which can be random or systematic. *Random errors* are characterized by the property that the mean of a sufficiently large number of independent measurements of the same object is arbitrary close to the true value. This does not hold for *systematic errors*, like those occurring when the tare is not compensated when weighing objects. Sometimes the term *precision* is used to describe the number significant figures, and the term *accuracy* to describe the distance to the real value (which is usually unknown, that is why we measure). The only defense against systematic errors is calibration, which is, of course, specific for the measurement. It bit more can be said about random errors, when we define them more exactly as the standard deviation. That is to say, we usually do not measure the size of a random error each time by repeating the measurement, but we use the value obtained from some prior calibration procedure (usually performed by the manufacturer of the equipment). The size of a random error can be expressed in terms of *absolute error*, having the same dimension as the measurement itself, or in terms of the

relative error, where we divide by the measured value, giving a dimensionless quantity.

Many measurements, like the surface area of a rectangular body, speed etc.) are compound, i.e. consist of a function of a number of other measurements. The following discussion shows how small independent random errors propagate in such compound measurements. For this purpose we linearize the function by a Taylor series approximation in the true value of the compound measurement. We explicitly neglect all terms involving powers of the error higher than 1, which is only reasonable if the error is small indeed. When f denotes the function of interest, $\underline{x}_1, \underline{x}_2, \dots$ the measurements, y_1, y_2, \dots the true values for the measurements, i.e. the expected values for x_1, x_2, \dots , so $\mathcal{E}\underline{x}_i = y_i$, and $\underline{d}_1, \underline{d}_2, \dots$ the deviations, i.e. $\underline{d}_i = \underline{x}_i - y_i$, we approximate $f(x_1, x_2, \dots)$ by

$$f(\underline{x}_1, \underline{x}_2, \dots) \simeq f(y_1, y_2, \dots) + \sum_i \underline{d}_i f'_i(y_1, y_2, \dots), \quad (4.1)$$

where f'_i represents the derivative of f with respect to x_i : $f'_i(y_1, y_2, \dots) = \frac{d}{dx_i} f(y_1, y_2, \dots)$. The first useful observation, when we apply the expectation operator to both sides, is that

$$\mathcal{E}f(\underline{x}_1, \underline{x}_2, \dots) \simeq \mathcal{E}f(y_1, y_2, \dots) + \sum_i \mathcal{E}\underline{d}_i f'_i(y_1, y_2, \dots) = f(y_1, y_2, \dots)$$

So the expected value of a function of a set of random variables is approximated by the function of the expectation of those random variables. (It should be stressed that this approximation is only close for really small deviations indeed). When we take the variance at both sides of (4.1), and apply the rules $\text{var } c = 0$ and $\text{var } (c\underline{x}) = c^2 \text{var } \underline{x}$ for a constant c , and $\text{var } (\underline{x}_1 + \underline{x}_2) = \text{var } \underline{x}_1 + \text{var } \underline{x}_2$ for independent \underline{x}_1 and \underline{x}_2 , we arrive at

$$\text{var } f(\underline{x}_1, \underline{x}_2, \dots) \simeq \text{var } f(y_1, y_2, \dots) + \text{var} \left(\sum_i \underline{d}_i f'_i(y_1, y_2, \dots) \right) \simeq (f'_i(y_1, y_2, \dots))^2 \sum_i \text{var } \underline{x}_i$$

The square of the relative error is thus given by

$$\frac{\text{var } f(\underline{x}_1, \underline{x}_2, \dots)}{f(y_1, y_2, \dots)^2} \simeq \left(\frac{d}{dx_i} \ln f(y_1, y_2, \dots) \right)^2 \sum_i \text{var } \underline{x}_i \quad (4.2)$$

From (4.2), we obtain the following squared relative errors which are of practical interest:

$$\frac{\text{var } (\underline{x}_1 + \underline{x}_2)}{(y_1 + y_2)^2} = \frac{\text{var } \underline{x}_1 + \text{var } \underline{x}_2}{(y_1 + y_2)^2} \quad (4.3)$$

$$\frac{\text{var } (\underline{x}_1 - \underline{x}_2)}{(y_1 - y_2)^2} = \frac{\text{var } \underline{x}_1 + \text{var } \underline{x}_2}{(y_1 - y_2)^2} \quad (4.4)$$

$$\frac{\text{var } (\underline{x}_1 \underline{x}_2)}{(y_1 y_2)^2} \simeq \frac{\text{var } \underline{x}_1}{y_1^2} + \frac{\text{var } \underline{x}_2}{y_2^2} \quad (4.5)$$

$$\frac{\text{var } (\underline{x}_1 / \underline{x}_2)}{(y_1 / y_2)^2} \simeq \frac{\text{var } \underline{x}_1}{y_1^2} + \frac{\text{var } \underline{x}_2}{y_2^2} \quad (4.6)$$

$$\frac{\text{var } \underline{x}^n}{y^{2n}} \simeq n^2 \frac{\text{var } \underline{x}}{y^2} \quad (4.7)$$

$$\frac{\text{var } n\underline{x}}{(ny)^2} = \frac{\text{var } \underline{x}}{y^2} \quad (4.8)$$

The relations (4.3), (4.4) and (4.8) also follow directly by applying the variance operator. They are exact rather than approximative. If the relative errors of the components of a compound measurement are known, (4.3)-(4.8) can be used to obtain the relative error of the compound measurement. As a rule, the last significant figure in the presentation of the (compound) measurement should be of the same order of magnitude as the error.

Example: If we measure that a daphnid of length 3.0 ± 0.2 mm, ingested $4.0 \pm 0.5 \cdot 10^5$ algal cells in 1.00 ± 0.01 h, and if we know that the ingestion rate is proportional to the squared length, than we find an ingestion rate for a 4 mm daphnid of $4.0 \cdot 10^5 \times 42/3.02 = 7.111 \cdot 10^5$ cells mm²/h. The relative error is $((0.5/4.0)^2 + (0.01/1.00)^2 + 4(0.2/3.0)^2)^{1/2} = 0.14$, so the absolute error is $7.11 \cdot 10^5 \times 0.14 = 1.0 \cdot 10^5$ cells.mm²/h, leading to the final presentation of $7.1 \pm 1.0 \cdot 10^5$ cells.mm²/h. Note that the largest relative error, that of the number of ingested cells in this case, dominates the error in the end result.

4.4 Smoothing and interpolation

If no strong preference for a particular model function exists, *cubic splines* can be used to interpolate between data points [?]. Cubic splines have parameters, called *knots* (or also nodes or joints), consisting of a sequence of point coordinates $\{x_i, y_i\}_{i=1}^n$. Between two neighbouring knots, the cubic spline is a 3-th degree polynomial; left of the first knot and right of the n -th knot, the cubic spline is a first degree polynomial. The coefficients of the $n - 1$ piecewise cubic polynomials and the 2 line segments are determined by the constraints that the cubic spline hits all knots, and that first and second derivatives of the polynomials left and right of each knot are equal. So one cannot see the knots in a graph of the spline; the $n + 1$ pieces are glued smoothly.

Given the knots $\{x_i, y_i\}_{i=1}^n$ for $n > 3$, and $d_i = x_{i+1} - x_i$, $\Delta_i = y_{i+1} - y_i$, the second derivatives at the knots, y''_2, \dots, y''_{n-1} for $y''_i = y''(x_i)$, can be found from

$$d_{i-1}y''_{i-1} + 2(d_{i-1} + d_i)y''_i + d_iy''_{i+1} = (\Delta_i/d_i - \Delta_{i-1}/d_{i-1})/6 \quad \text{for } i = 2, \dots, n - 1$$

while $y''_1 = y''_n = 0$. We will need $\Delta''_i = y''_{i+1} - y''_i$. The first derivatives at the knots, y'_2, \dots, y'_{n-1} for $y'_i = y'(x_i)$, are

$$y'_i = \Delta_i/d_i - (2y''_i + y''_{i+1})d_i/6 \quad \text{for } i = 1, \dots, n - 1; \quad y'_n = \Delta_{n-1}/d_{n-1} + (y''_{n-1} + 2y''_n)d_{n-1}/6$$

We will also need $\Delta'_i = y'_{i+1} - y'_i$. The cubic spline is now given by

$$y(x) = \begin{cases} y_1 - (x_1 - x)y'_1 & \text{for } x \leq x_1 \\ y_i + (x - x_i)\Delta_i/d_i - (x - x_i)(x_{i+1} - x)(y''_i + y''(x) + y''_{i+1})/6 & \text{for } x_i < x \leq x_{i+1}, i = 1, \dots, n - 1 \\ y_n + (x - x_n)y'_n & \text{for } x > x_n \end{cases}$$

The first derivative is

$$y'(x) = \begin{cases} \Delta_1/d_1 - d_1y''_2/6 & \text{for } x \leq x_1 \\ \Delta_i/d_i + (2x - x_i - x_{i+1})(y''_i + y''(x) + y''_{i+1})/6 & \text{for } x_i < x \leq x_{i+1}, i = 1, \dots, n - 1 \\ \Delta_{n-1}/d_{n-1} + d_{n-1}y''_{n-1}/6 & \text{for } x > x_n \end{cases}$$

The second derivative is

$$y''(x) = \begin{cases} 0 & \text{for } x \leq x_1 \\ y_i'' + (x - x_i)\Delta_i''/d_i & \text{for } x_i < x \leq x_{i+1}, i = 1, \dots, n-1 \\ 0 & \text{for } x > x_n \end{cases}$$

The third derivative is

$$y'''(x) = \Delta_i''/d_i \quad \text{for } x_i < x \leq x_{i+1}, i = 1, \dots, n-1 \quad \text{and } y'''(x) = 0 \text{ for } x < x_1 \text{ and } x > x_n$$

The integral is $\int_a^b y(x), dx =$

$$\begin{cases} (b-a)y_1 - (b-a)(x_1 - a/2 - b/2)y_1' & \text{for } a \leq b \leq x_1 \\ (b-a)y_i + (b-a)(a/2 + b/2 - x_i)\Delta_i/d_i + \\ \quad + \frac{1}{6}(x_i x_{i+1}(b-a) - (x_{i+1} + x_i)(b^2 - a^2)/2 + (b^3 - a^3)/3)(2y_i'' + y_{i+1}'' - x_i\Delta_i''/d_i) + \\ \quad + \frac{1}{6}(x_i x_{i+1}(b^2 - a^2)/2 - (x_{i+1} + x_i)(b^3 - a^3)/3 + (b^4 - a^4)/4)\Delta_i''/d_i \\ \quad \text{for } x_1 \leq a \leq b \leq x_{i+1}, \quad i = 1, \dots, n-1 \\ (b-a)y_n - (b-a)(x_n - a/2 - b/2)y_n' & \text{for } x_n \leq a \leq b \end{cases}$$

If the knots are all data points, the spline is called an interpolation spline. Scatter in the data can easily lead to erratic behaviour of the interpolating spline. In practice fewer knots are chosen, i.e. the values for the independent variable are selected, and those for the dependent variable are chosen such that the sum of squared distances between the data points and the spline is minimized. Such a spline is called a smoothing spline. The fewer the number of knots, the smoother the spline, but 4 is the minimum. It is strongly recommended to check spline fittings graphically.

The advantage of splines is that the shape of the function in one interval of the domain, hardly affects the shape in another interval; this in contrast to a single high-degree polynomial, for instance.

Important applications, apart from interpolation, are in numerical differentiation, in integration and in root finding, if only data points are available; the data need not be measured values, but can also be calculated values. Notice that the second derivative of the cubic spline is still continuous, but the third derivative does not exist at the knots.

What splines do for functions, Bézier curves do for isoclines. These are functions of the type

$$p(u) = \sum_{i=0}^n p_i \binom{n}{i} u^i (1-u)^{n-i}$$

where p_i is a *control point*, and n some chosen number, frequently 3. While u walks from 0 to 1, the curve $p(u)$ connects point $p(0) = p_0$ with $p(1) = p_n$. The curve does not hit intermediary control points; at p_0 , the curve is tangent to the line segment (p_0, p_1) , and at p_n the curve is tangent to the line segment (p_{n-1}, p_n) . If p_0 and p_n coincide, the curve is closed; a circle results if p_1, \dots, p_4 are on the corners of a square, and $p_0 = p_5$ are in the middle of an edge. The point p_0 of the new segment $i+1$, called p_0^{i+1} , will generally coincide with point p_n^i of the last segment. If $p_{n-1}^i, p_n^i = p_0^{i+1}, p_1^{i+1}$ are on the same line, the joint between two adjacent Bézier curves is smooth. A problem of Bézier curves is that the whole segment changes in shape, if one control point is changed. *B-splines* are used to solve these problems [?].




Figure 4.1: The 20 data points are indicated by red +, the 6 knots with blue +, the smoothing cubic spline in green, the first derivative in blue, the second one in magenta, the third one in red.

4.5 Testing hypotheses

A statistical test usually takes to form of

- formulation of a *null* (H_0) and an *alternative* (H_1) hypothesis about the value of a parameter, e.g. $H_0: p = p_0$ and $H_1: p = p_1$ (a so-called *simple* alternative) or e.g. $H_1: p \neq p_0$ (a *composite* alternative)
- definition of a test statistic, which is some appropriate function of the observations
- derivation of the distribution of the test statistic under the null hypothesis
- rejection of the null hypothesis if the survivor function under H_0 at the value of the test statistic is less than some preselected value: the *significance level*.

Two types of errors can occur in this decision scheme

- *Type 1 error*: H_0 is rejected, while it is true. This occurs with a probability that is equal to the significance level.
- *Type 2 error*: H_0 is accepted, while H_1 is true. This occurs with a probability that is equal to the 1 minus the *power* of the test. The distribution of the test statistic under H_1 must be derived to calculate the power.

The *Neyman-Pearson Theorem* states that, if x_1, \dots, x_n represent n independent trials from some p.d.f., the best (i.e. uniformly most powerful) choice for a test statistic for simple alternatives is the likelihood ratio (see below).

Example: a test on the value of the binomial probability is as follows. Suppose that we have a trial x from the binomial probability distribution with parameters n and p , and want to test the null hypothesis $H_0: p = 0.2$ against the alternative $H_1: p = 0.4$. We choose the significance level α and obtain the smallest value for m for which $\sum_{i=m}^n \Pr\{\underline{x} = i | p = 0.2\} \leq \alpha$ holds. We reject H_0 if $x > m$. It is in this case usually not possible to test exactly with significance level α because x is an integer. The power of the test is $1 - \sum_{i=m}^n \Pr\{\underline{x} = i | p = 0.4\}$.

4.6 Likelihood functions

Consider a random sample $\underline{x}_1, \underline{x}_2, \dots, \underline{x}_n$ of some discrete random variable that has probability distribution $\Pr\{\underline{x} = x; \theta\}$, where θ represents the vector of parameters. If the samples are independent, the simultaneous probability of $\underline{x}_1, \underline{x}_2, \dots, \underline{x}_n$ is $\prod_{i=1}^n \Pr\{\underline{x}_i = x_i; \theta\}$. This simultaneous probability may be regarded as a function of θ . When so regarded, it is called the *likelihood function* L of the random sample, and we write

$$L(\theta; x_1, x_2, \dots, x_n) = \prod_{i=1}^n \Pr\{\underline{x}_i = x_i; \theta\}$$

The value of θ that maximizes L is called the *maximum likelihood estimate* (MLE) of θ and it will be denoted as $\hat{\theta}$ and is a number. If we replace the observed values x_i by the random variables \underline{x}_i in the MLE, we have an *maximum likelihood estimator*, $\hat{\theta}$, which is a random variable. Be aware of the problem that likelihood functions usually have several local maxima, while the MLE only relates to the global maximum.

Example with real data: Suppose we have the following 10 observations from a Poisson variable with unknown parameter λ : 4, 0, 2, 5, 3, 2, 4, 1, 7, 3. Which value of λ fits best to the data? According to the maximum likelihood principle we have to maximize the product of the 10 Poisson probabilities

$$\begin{aligned} L(\lambda; 4, 0, \dots, 3) &= f(4; \lambda)f(0; \lambda) \dots f(3; \lambda) \\ &= \frac{\lambda^4}{4!}e^{-\lambda} \frac{\lambda^0}{0!}e^{-\lambda} \dots \frac{\lambda^3}{3!}e^{-\lambda} \\ &= \frac{\lambda^{31}}{4!0! \dots 3!}e^{-10\lambda} \end{aligned}$$

This function can be maximized by setting the first derivative of L , with respect to λ , equal to zero. Note however, that each of the functions L and $\ln(L)$ (denoted as ℓ) is a maximum for the same value of λ . It appears to be easier to work with ℓ instead of L .

$$\begin{aligned} \frac{d\ell}{d\lambda} &= 0 \\ \ell(\lambda) &= 31 \ln(\lambda) - \ln(4!0! \dots 3!) - 10\lambda \\ \frac{d\ell}{d\lambda} &= \frac{31}{\lambda} - 10 = 0 \leftrightarrow \lambda = \frac{31}{10} = 3.1. \end{aligned}$$

So the MLE $\hat{\lambda} = 3.1$, the average of the data. From this simple example we can derive a more general result: if we have a sample from a Poisson variable, the maximum likelihood estimator of the parameter λ is given by

$$\hat{\lambda} = \frac{\sum \underline{x}_i}{n} = \bar{\underline{x}}$$

For n independent samples from a continuous random variable, we can likewise define the likelihood function

$$L(\theta; x_1, x_2, \dots, x_n) = \prod_{i=1}^n f_{\underline{x}_i}(x_i; \theta)$$

with $\dim(L) = \dim(x)^{-n}$. An example for the exponentially distributed random variable with unknown parameter λ is as follows. Let $\underline{x}_1, \underline{x}_2, \dots, \underline{x}_n$ be such a sample of size n . The likelihood L is now given by the simultaneous p.d.f.

$$\begin{aligned} L(\lambda; x_1, \dots, x_n) &= f(x_1; \lambda)f(x_2; \lambda) \dots f(x_n; \lambda) \\ &= \lambda e^{-\lambda x_1} \lambda e^{-\lambda x_2} \dots \lambda e^{-\lambda x_n} \\ &= \lambda^n e^{-\lambda \sum x_i} \end{aligned}$$

As in the previous example we take the logarithm of L .

$$\ell(\lambda) = n \ln(\lambda) - \lambda \sum x_i$$

From this the maximum likelihood estimator can easily be found to be

$$\hat{\lambda} = \frac{n}{\sum x_i} = \frac{1}{\bar{x}}$$

Example with two parameters: the normal distribution. Suppose we have a sample from a normal distribution $N(\mu, \sigma^2)$. The log likelihood ℓ is then given by

$$\ell(\mu, \sigma^2) = -\frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln(\sigma^2) - \frac{\sum_i (x_i - \mu)^2}{2\sigma^2} \quad (4.9)$$

To find the maximum we take the two partial derivatives with respect to μ and σ^2 and equate these to zero.

$$\begin{aligned} \frac{\partial \ell}{\partial \mu} &= \frac{1}{\sigma^2} \sum (x_i - \mu) \\ \frac{\partial \ell}{\partial \sigma^2} &= -\frac{n}{2\sigma^2} + \frac{1}{2\sigma^4} \sum (x_i - \mu)^2 \end{aligned}$$

This leads to the following MLE's for μ and σ^2

$$\hat{\mu} = \bar{x} \quad \text{and} \quad \hat{\sigma}^2 = \frac{1}{n} \sum (x_i - \bar{x})^2 \quad (4.10)$$

4.7 Large sample properties of ML estimators

The estimator $\hat{\theta}$ is said to be an *unbiased* estimator for θ if $\mathcal{E}\hat{\theta} = \theta$. ML estimators are generally only asymptotically unbiased, i.e. for increasingly large sample sizes.

The ML estimator for the Poisson parameter λ has nice properties: it's obvious that $\mathcal{E}(\bar{x}) = \lambda$ and $\text{var}(\bar{x}) = \lambda/n$, so $\mathcal{E}(\hat{\lambda}) = \lambda$ and $\text{var}(\hat{\lambda}) = \lambda/n$. The first property means that the maximum likelihood estimator for the Poisson parameter is unbiased. We also know the distribution of the estimator: $\sum x_i \sim \text{Poisson}(n\lambda)$, so $\hat{\lambda} \sim \frac{1}{n} \text{Poisson}(n\lambda)$.

The ML estimator for the exponential parameter λ has not the same nice properties as in the previous example. Here it is not straightforward how to calculate the expectation

of $\hat{\lambda}$. It appears that $\mathcal{E}(\frac{1}{\hat{x}}) = \frac{n}{n-1}\lambda$, i.e. $\hat{\lambda}$ is biased. But for large n the bias disappears. We say that $\hat{\lambda}$ is asymptotically unbiased. This appears to be a general result.

It can be shown that ML estimators have, asymptotically, *minimum variance*, i.e. no other estimators can be devised that have a smaller variance than ML estimators for increasingly large sample sizes.

The ML theory also applies to a vector of parameters. Under general regularity conditions it can be proved that a maximum likelihood estimator $\hat{\underline{\theta}}$ follows asymptotically (i.e. for large values of the sample size n) a normal distribution:

$$\hat{\underline{\theta}} \sim N(\underline{\theta}, \Sigma(\underline{\theta})) \quad \text{where} \quad \Sigma(\underline{\theta}) = \left[-\mathcal{E} \left(\frac{d^2 \ell}{d\theta d\theta'} \right) \right]^{-1} = \left[\mathcal{E} \left(\frac{d\ell}{d\theta} \frac{d\ell}{d\theta'} \right) \right]^{-1}$$

In the case of a Poisson distribution this leads to

$$\hat{\lambda} \sim N\left(\lambda, \frac{\lambda}{n}\right) \quad \text{for large values of } n$$

For the exponential distribution we get

$$\hat{\lambda} \sim N\left(\lambda, \frac{\lambda^2}{n}\right) \quad \text{for large values of } n$$

The asymptotic distributions can be used to construct *confidence intervals* for the parameter; it is said to be of level α with boundaries θ_0 and θ_1 if $\Pr\{\theta_0 \leq \hat{\theta} \leq \theta_1\} = \alpha$. Notice that such an interval is not unique. A different way to get confidence intervals is by the use of profile likelihoods.

4.8 Likelihood ratio principle

The likelihood function can be used to construct parameter estimators, but it can also be used to test hypotheses about those parameters. The idea is as follows: we take a sample from a distribution with unknown parameter vector $\theta = (\theta_1, \theta_2, \dots, \theta_k)$ that can take values in a set Ω . We want to test a hypothesis, for instance $\theta_1 = 0$ or more generally $\theta \in \Omega_0$.

Now we maximize the likelihood in two ways: first by constraining ourselves to Ω_0 , next without any constraint. Then we look at LR , the ratio of the two likelihoods:

$$LR = \frac{\max_{\theta \in \Omega_0} L(\theta)}{\max_{\theta \in \Omega} L(\theta)}$$

If LR is too small we have to reject the null hypothesis. What we mean by 'too small' is determined by the distribution of LR in the usual way: a critical value c is given by $P(LR < c | H_0) = \alpha$ where α is the significance level (usually 0.05). Note that we can also work with the logarithm of LR , that is the difference between the logarithms: $\max_{\theta \in \Omega_0} \ell(\theta) - \max_{\theta \in \Omega} \ell(\theta)$. Because it is always negative, we usually work with the *deviance* $-2 \ln LR \equiv \mathcal{L} = 2(\max_{\theta \in \Omega} \ell(\theta) - \max_{\theta \in \Omega_0} \ell(\theta))$, which is always positive.

This idea has been worked out for a lot of standard experiments: the well known t -test, ANOVA and linear regression tests are all examples of likelihood ratio tests, based on the normal distribution. The binomial test for binary data (0/1, yes/no, blue/green) is also based on the same principle. These tests are exact: they do not assume large samples.

There is a large sample approach that is generally applicable: it can be proved that asymptotically

$$\underline{\mathcal{L}} \sim \chi_{\nu}^2$$

where ν is the difference in the number of parameters to be estimated, between maximization over Ω and Ω_0 . Graphically this means that we first approximate the log likelihood ratio around the MLE by the tangent parabola:

$$\mathcal{L}(\theta) \simeq -(\theta - \hat{\theta})^2 \frac{d^2}{d\theta^2} \ell(\hat{\theta})$$

The boundary values of the $100(1 - \alpha)\%$ confidence region are found by the intersection of this parabola with the line $\mathcal{L}(\theta) = \chi_{[\nu; 1-\alpha]}^2$, which gives

$$\hat{\theta} - \sqrt{-\chi_{[\nu; 1-\alpha]}^2 / \frac{d^2}{d\theta^2} \ell(\hat{\theta})} < \theta < \hat{\theta} + \sqrt{-\chi_{[\nu; 1-\alpha]}^2 / \frac{d^2}{d\theta^2} \ell(\hat{\theta})}$$

A simple example: suppose we have data x_1, x_2, \dots, x_n from an exponential distribution and $\Omega = (0, \infty)$. The ln-likelihood is given by $\ell(\lambda) = n(\ln(\lambda) - \lambda\bar{x})$. The maximum value of ℓ over Ω is reached for $\lambda = \hat{\lambda} = 1/\bar{x}$ and amounts to $\ell(\hat{\lambda}) = -n(\ln(\bar{x}) + 1)$, so $\frac{d^2}{d\lambda^2} \ell(\hat{\lambda}) = -n\hat{\lambda}^{-2} = -n\bar{x}^2$ and $\mathcal{L}(\lambda) = 2n(\lambda\bar{x} - 1 - \ln(\lambda\bar{x}))$. The tangent parabola approximation is $\mathcal{L}(\lambda) \simeq n(\lambda\bar{x} - 1)^2$, which gives the $100(1 - \alpha)\%$ confidence interval

$$\bar{x} \left(1 - \sqrt{\chi_{[1; 1-\alpha]}^2 / n}\right) < \lambda < \bar{x} \left(1 + \sqrt{\chi_{[1; 1-\alpha]}^2 / n}\right)$$

see Figure 4.2. If we want to test the hypothesis $\lambda = 0.5$, we have $\Omega_0 = \{0.5\}$. Maximization of L over Ω_0 is simple: it is $L(0.5)$. This leads to the following test rule: reject $H_0 : \lambda = 0.5$ if

$$2n(0.5\bar{x} - 1 - \ln(0.5\bar{x})) > \chi_{[1; 1-\alpha]}^2$$

4.8.1 Likelihood based confidence region

We can use the previous approach to construct a confidence region for a parameter vector $\theta = (\theta_1, \theta_2, \dots, \theta_k)$. The region consists of all parameter vector values θ_0 that do not lead to rejection of the hypothesis $H_0 : \theta = \theta_0$. This leads to the $100(1 - \alpha)\%$ confidence region

$$\left\{ \theta, 2 \ln \left(\frac{L(\hat{\theta})}{L(\theta)} \right) < \chi_{[\nu; 1-\alpha]}^2 \right\}$$




Figure 4.2: The $-2 \ln$ likelihood ratio, \mathcal{L} , as a function on its parameter λ , together with the tangent parabolic approximation for 5 samples from an exponential distribution. The range for which the functions are below the indicated horizontal line represents the 95 % confidence interval. \mathcal{L} is here proportional to the sample size; an increase of the number of samples has the same effect as lowering the horizontal line. Close to the MLE, the parabolic approximation is very good, but for small samples, it should not be used.

4.8.2 Profile likelihood

The confidence region defined above becomes difficult to handle or to communicate in case of multi-parameter models. In most case we are interested in one parameter at a time, say in θ_1 . Therefore we calculate the profile likelihood $L_p(\theta_1)$ given by

$$L_p(\theta_1) = \max_{\theta_2, \dots, \theta_k} L(\theta_1, \theta_2, \dots, \theta_k)$$

In case of a one-parameter model, the boundary values of the $100(1-\alpha)\%$ confidence region are found by the intersection of the $-2 \ln$ likelihood ratio with the line $\mathcal{L}(\theta) = \chi_{[p; 1-\alpha]}^2$. In the section on likelihood based confidence region we used the tangent parabola, rather than the likelihood ratio itself. The boundary values of the confidence region are no longer equidistant from the MLE. In the case of the exponential distribution the boundary values of the $100(1-\alpha)\%$ confidence interval have to be found numerically from $\mathcal{L}(\lambda) = 2n(\lambda\bar{x} - 1 - \ln(\lambda\bar{x})) = \chi_{[1; 1-\alpha]}^2$, see Figure 4.2. The purpose is to demonstrate the general applicability of the profile likelihood method; the fact that the exact confidence interval can be obtained analytically in this particular case is of no relevance.

In case of a two-parameter model the likelihood function can be seen as a mountain landscape. The profile likelihood is then the skyline of this landscape in one dimension. From this one-dimensional function we can calculate a $100(1-\alpha)\%$ confidence interval for θ_1 given by

$$\left\{ \theta_1, 2 \ln \left(\frac{L(\hat{\theta})}{L_p(\theta_1)} \right) < \chi_{[1; 1-\alpha]}^2 \right\}$$

The confidence region is based on large samples. Profile likelihood ratio's tend to give more correct confidence intervals for small samples than approximations by tangent parabolas. The practical problem in the application of profile likelihoods is that for each value of the parameter under consideration, the MLE's for the other parameters have to be obtained, which might be computationally intensive.

In case of a likelihood function with more than one (local) maximum this definition of a confidence region may lead to more than one disjunct intervals. In general it would be better to use the term confidence set instead of confidence interval.

4.9 Regression

Regression models are deterministic models for dependent variables with additive “scatter” frequently have the form

$$y_i(x) \sim N(f(x_i; \theta), \sigma^2)$$

where f is some specified function of an independent variable with parameter vector θ , and σ a (constant) scatter parameter. The ln likelihood function is

$$\ell(\theta, \sigma^2) = -\frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln(\sigma^2) - \frac{\sum_i (y_i - f(x_i; \theta))^2}{2\sigma^2}$$

The maximization of this likelihood function as a function of the parameters amounts to a minimization of the sum of squared deviations between the observed values y_i and the (deterministic) model predictions $f(x_i)$. The least squared deviation criterion for parameter estimation is thus a special case of the ML criterion. The ML estimator for σ^2 is for $f_i = f(x_i; \theta)$

$$\hat{\sigma}^2 = \frac{1}{n} \sum_i (y_i - f_i)^2.$$

If we want to test the hypothesis that p parameters have specified values, based on n observations, the deviance amounts to $\mathcal{L} = n \log \hat{\sigma}_1^2 / \hat{\sigma}_0^2$, where $\hat{\sigma}_1^2$ is the estimated variance if all parameters are estimated, while $\hat{\sigma}_0^2$ is that given the values of the p parameters. Under the null hypothesis, \mathcal{L} is χ^2 distributed with p degrees of freedom.

4.9.1 Constant variation coefficient

Biological data tend to have a constant variation coefficient (vc), so $\sigma = v_c f(x; \theta)$ and $f(x; \theta) > 0$, rather than a constant variance. For normally distributed dependent variables ln likelihood function now amounts to

$$\ell(\theta, v_c) = -\frac{n}{2} \ln(2\pi) - n \ln v_c - \sum_i \ln f(x_i; \theta) - \frac{1}{2v_c^2} \sum_i (y_i / f(x_i; \theta) - 1)^2$$

This leads for $f_i = f(x_i; \theta)$ to

$$\hat{v}_c^2 = \frac{1}{n} \sum_i (y_i / f_i - 1)^2$$

The ML estimates minimize $\log \hat{\sigma}$ in the constant-sd model, and $\log \hat{v}_c + \frac{1}{n} \sum_i \log f_i$ in the proportional-sd model. The first and second derivatives of the ln likelihood function with respect to the parameters and an estimate for the variance of the ML estimator are

$$\begin{aligned} \frac{d}{d\theta} \ell &= -\sum_i \frac{\frac{d}{d\theta} f_i}{f_i} + \frac{1}{v_c^2} \sum_i \left(\frac{y_i}{f_i} - 1 \right) \frac{y_i \frac{d}{d\theta} f_i}{f_i^2} = 0 \quad \text{for } \theta = \hat{\theta} \\ \frac{d^2}{d\theta d\theta^T} \ell &= -\frac{1}{v_c^2} \sum_i \frac{y_i \left(\frac{d}{d\theta} f_i \right) \left(\frac{d}{d\theta^T} f_i \right)}{f_i^3} - \sum_i \left(\frac{\frac{d^2}{d\theta d\theta^T} f_i}{f_i} - \frac{\left(\frac{d}{d\theta} f_i \right) \left(\frac{d}{d\theta^T} f_i \right)}{f_i^2} \right) \left(1 - \frac{1}{v_c^2} \frac{y_i}{f_i} \left(\frac{y_i}{f_i} - 1 \right) \right) \end{aligned}$$




Figure 4.3: Data with curves $y(x) = y_\infty - (y_\infty - y_0) \exp(-rx)$, with ML estimated parameters, based on the assumption of normally distributed scatter with constant standard deviation (in red: $\hat{y}_0 = 0.023$, $\hat{y}_\infty = 2.98$, $\hat{r} = 0.695$) and with standard deviation proportional to the mean (in green: $\hat{y}_0 = 0.003$, $\hat{y}_\infty = 3.01$, $\hat{r} = 0.686$). The small difference is rather typical, but exceptions can occur.

$$\begin{aligned} &\simeq -\frac{1}{v_c^2} \sum_i \frac{y_i \left(\frac{d}{d\theta} f_i\right) \left(\frac{d}{d\theta^T} f_i\right)}{f_i^3} \\ \text{var}(\hat{\theta}) &= -\left(\mathcal{E} \frac{d^2}{d\theta d\theta^T} \ell\right)^{-1} \simeq v_c^2 \left(\sum_i \frac{\left(\frac{d}{d\theta} f_i\right) \left(\frac{d}{d\theta^T} f_i\right)}{f_i^2}\right)^{-1} \end{aligned}$$

Scoring iteration scheme to find $\hat{\theta}$ from θ_0 where f_i is evaluated at θ_j

$$\theta_{j+1} = \theta_j + \text{var}(\hat{\theta}) \frac{d}{d\theta} \ell = \theta_j + \left(\sum_i \frac{\left(\frac{d}{d\theta} f_i\right) \left(\frac{d}{d\theta^T} f_i\right)}{f_i^2}\right)^{-1} \sum_i \frac{d}{d\theta} f_i \left(\left(\frac{y_i}{f_i} - 1\right) \frac{y_i}{f_i} - \hat{v}_c^2\right)$$

4.10 Composite likelihoods

In the practice of statistical analyses of biological data, it frequently happens that not a single, but several variables have been measured and that particular parameters occur in more than one data set. A very much related problem is the case that the same variable has been measured in different experiments, while some parameters have the same value and others (for instance parameters that relate to “scatter”) have different values for the different experiments. As long as the data are mutually independent the method of estimating the parameters (point and interval estimates) is straightforward within the context of the likelihood principle. We illustrate this with a simple example.

Suppose we have data sets $\{x_i, y_i\}_{i=1}^n$ and $\{v_i, w_i\}_{i=1}^m$. We assume that the random variable \underline{y} is normally distributed with mean ax and variance σ_y^2 , while the random variable \underline{w} is also normally distributed with mean av and variance σ_w^2 . The ln likelihood function for this case is

$$\ell(a, \sigma_y, \sigma_w) = -\frac{n}{2} \ln 2\pi\sigma_y^2 - \frac{1}{2\sigma_y^2} \sum_{i=1}^n (y_i - ax_i)^2 - \frac{m}{2} \ln 2\pi\sigma_w^2 - \frac{1}{2\sigma_w^2} \sum_{i=1}^m (w_i - av_i)^2$$

The maximization of the ln likelihood function, as a function of the parameter a amounts the a minimization of the weighted sum of squared deviations between measured and expected values for y_i and w_i (terms 2 and 4 in the right-hand side of the formula). The

weight coefficients are inversely proportional to the variances. The resulting MLE for the parameter a turns out to be

$$\hat{a} = \left(\hat{\sigma}_y^{-2} \sum_{i=1}^n x_i y_i + \hat{\sigma}_w^{-2} \sum_{i=1}^m v_i w_i \right) \left(\hat{\sigma}_y^{-2} \sum_{i=1}^n x_i^2 + \hat{\sigma}_w^{-2} \sum_{i=1}^m v_i^2 \right)^{-1}$$

while the variances are estimated by

$$\hat{\sigma}_y^2 = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{a} x_i)^2 \quad \text{and} \quad \hat{\sigma}_w^2 = \frac{1}{m} \sum_{i=1}^m (w_i - \hat{a} v_i)^2$$

Notice that $\hat{\sigma}^2 = \hat{\sigma}_y^2$; this is just a special case of a general property of MLE's: the MLE of a (parameter-free) function of a parameter is the function of the MLE of the parameter. Notice also that the estimates for \hat{a} , $\hat{\sigma}_y$ and $\hat{\sigma}_w$ are given implicitly only; explicit estimates can only be obtained numerically.

4.11 Parameter identifiability

A parameter is said to be *unidentifiable* if it cannot be estimated from a given data set. It is theoretically unidentifiable if it can never be estimated, no matter how extensive the data set is. The parameters a and x_0 in the model $y(x) = a(x/x_0)^b + \epsilon$, for example, cannot be estimated from a data set $\{x_i, y_i\}_{i=1}^n$. The solution to this problem is re-parametrization, using less parameters. Here we have $y(x) = a'x^b$, where $a' = ax_0^{-b}$ (although this here gives a dimension problem, which strangely does not stop many workers from using this model.).

A parameter set is practically unidentifiable if it cannot be estimated from the given data set, but this problem might go away for other data sets. The parameters y_m and k in the model $y(x) = \frac{y_m x}{k+x} + \epsilon$ can hardly be estimated from data set $\{x_i, y_i\}_{i=1}^n$, if $\max_i x_i < k$. The parameters y_m and k are strongly negatively correlated. The solution to this problem is either extending the data set or re-parametrization with less parameters. If $x < k$ we have $y(x) \simeq y'_m x + \epsilon$, where $y'_m = y_m/k$.

The properties of parameter estimators depend on the way the parameters are introduced. In the regression of y on x , the estimators for parameters a and b in the relationship $y = x^2(a + bx)$ are strongly negatively correlated when in the observations $\{x_i, y_i\}_{i=1}^n$, all $x_i > 0$; the mathematically totally equivalent relationship $y = x^2(c + b(x - \sum_i x_i^3 / \sum_i x_i^2))$ suffers much less from this problem. Replacement of the original parameters by appropriately chosen compound parameters can also reduce correlations between parameter estimates.

4.12 Monte Carlo techniques

The possibilities to evaluate the properties of MLEs for small samples are very limited; likelihood functions are even difficult to construct for interacting sub-processes. The construction of such functions are only feasible if we deal with *independent* trials from some




Figure 4.4: Left: The empirical survivor function of the ML estimator based on 2, 5 and 10 random trials from an exponential distribution with parameter 1. Right: the corresponding empirical confidence level, against the likelihood-based confidence level. All curves were generated from 1000 MLE's.

distribution. The only method that is frequently left is the Monte Carlo method, where computer simulations are used to evaluate the propagation of effects of stochasticity.

The basic tool is the *random generator*, which is a recurrent deterministic algorithm; a next trial from an uncorrelated approximately homogeneously distributed random variable is obtained from a previous one, after initialization with a *seed*. These numbers are subsequently transformed to produce random trials from other distributions and used to simulate the process under consideration. This has to be done many times to evaluate the role of stochasticity but that does not need to be a problem.

As an example of application of Monte Carlo methods, we check the validity of the likelihood-based confidence interval for the ML estimator for the parameter of the exponential distribution. The procedure is to choose a confidence level, obtain the likelihood-based confidence interval, and calculate the fraction of the MLEs that are in this interval. We repeat this for many choices of confidence levels. Figure 4.4 shows that even for 2 random trials, the likelihood-based confidence interval is close to correct. (The tail probabilities are in this case not exactly equal, however.)

Chapter 5

Notation

The following conventions are used:

- summation: $\sum_{i=1}^n x_i = x_1 + x_2 + \cdots + x_n$. If the summation is over all indices, this is summarized as $\sum_i x_i$. The double summation $\sum_{j=1}^m \sum_{i=1}^n x_{ij}$ is summarized as $\sum_{ij} x_{ij}$. The notation $\sum_{i \neq j} x_{ij}$ means: summation over all indices i and j such that $i \neq j$, so the elements x_{ii} are excluded. The notation $\sum_{i < j} x_{ij}$ means: summation over all indices i and j such that $i < j$.
- product: $\prod_{i=1}^n x_i = x_1 x_2 \cdots x_n$. Indices handling is similar to summation.
- integration: $\int_{x=a}^b f(x) dx$ or $\int_a^b f(x) dx$, where x runs from a till b . If the integration is over the whole domain, it is summarized as $\int_x f(x) dx$. The notation $\int f(x) dx$, however, means the indefinite integral.
- capitals for matrices
- τ for transposition
- \prime and $\frac{d}{dx}$ for differentiation
- diag for diagonal elements
- $||$ for determinant
- $|||$ for length
- underline for random variables
- Pr for probability
- var for variance
- cov for covariance
- cor for correlation coefficient

- vc for variation coefficient
- sd for standard deviation
- \mathcal{E} for expectation operator
- \dim for dimension

Index

- accuracy, 51
- antiderivative, 26
- approximation, 24
- Arrhenius, 6
- asymptote, 21
- attractor
 - cyclic, 45
 - point, 45
 - strange, 45

- base, 20
- bias, 57
- bifurcation, 45
- biology
 - experimental, 3
 - theoretical, 3

- chaos, 45
- circle, 28
- coefficient
 - correlation, 34
 - variation, 34
- coherence, 7
- complement, 17
- conjugate, 19
- consistency, 3
- constraint, 25
- continuation, 46
- contour, *see* isocline
- convergence, 22
- cosecant, 20
- cosine, 20, 33
- cotangent, 20
- covariance, 34
- curve
 - Bezier, 54

- definite
 - negative, 32
 - positive, 32
- definition
 - implicit, 20
 - recursive, 22
- degree, 19

- dependent
 - linearly, 30
 - stochastically, 34
- derivative
 - directional, 24
 - partial, 24
 - second, 23
- design
 - experimental, 10
- determinant, 25, 27, 29
- deviance, 58
- deviation
 - standard, 34
- diagram
 - bifurcation, 47
 - operating, 47
 - Venn, 17
- differentiation
 - implicit, 23
 - implicit partial, 24
 - numerical, 38
- dimension, 5
- disjoint, 17
- distribution
 - Bernoulli, 35
 - beta, 36
 - binomial, 35
 - Chi-square, 36
 - exponential, 36
 - frequency, 33
 - gamma, 36
 - geometric, 35
 - homogeneous, 36
 - marginal, 37
 - multinomial, 35
 - negative binomial, 35
 - normal, 36
 - Poisson, 35
 - probability, 33
 - truncated, 37
- divergence, 22
- domain, 18

- elasticity, 22

- element, 28
- ellipse, 28
- empirical cycle, 1
- equation
 - differential, 43
 - stiff differential, 38
- error
 - absolute, 51
 - random, 51
 - relative, 52
 - systematic, 51
 - type 1, 55
 - type 2, 55
- estimate
 - maximum likelihood, 56
- estimator
 - maximum likelihood, 56
- expectation, 34
- factorial, 18
- falsification, 3
- feedback, 45
- figure
 - significant, 51
- form
 - bilinear, 32
 - normal, 46
 - quadratic, 32
- function, 19
 - algebraic, 20
 - allometric, 6
 - beta, 26
 - composition, 20
 - continuous, 21
 - correlation, 41
 - differentiable, 22
 - distribution, 34
 - even, 19
 - exponential, 26
 - gamma, 26
 - homogeneous, 19
 - inner, 20
 - inverse, 20
 - likelihood, 56
 - logarithmic, 26
 - monotonous, 19
 - odd, 19
 - outer, 20
 - polynomial, 19
 - probability density, 33
 - rational, 19
 - Riemann zeta, 22
 - spline, 53
 - survivor, 34
 - transcendental, 20
- generator
 - random, 64
- gradient, 24
- hyperbola, 28
- hypothesis
 - alternative, 55
 - null, 55
- idempotent, 29
- image, 18
- input, 43
- integer, 19
- integral
 - convolution, 37
- integrand, 26
- integration
 - definite, 26
 - indefinite, 26
 - numerical, 38
 - partial, 27
- intensity, 43
- interpolation, 53
- intersection, 17
- interval, 18
 - closed, 18
 - confidence, 58
 - half-open, 18
 - open, 18
 - unbounded, 18
- isocline, 28
- iteration
 - Newton Raphson, 39
- Jacobian, 27
- knot, 53
- level
 - significance, 55
- limit, 21
 - left-hand, 21
 - right-hand, 21
- line
 - tangent, 24
 - vertical tangent, 21
- logic
 - deductive, 14
 - inductive, 14
 - predicate, 15
 - propositional, 14

- map, 41
- mapping, *see* operator
- Markov
 - chain, 42
 - process, 42
- matrix, 28
 - diagonal, 29
 - generalized inverse, 31
 - Hessian, 24
 - identity, 29
 - inverse, 30
 - Jacobian, 27
 - nonsingular, 30
 - rotation, 33
 - singular, 30
 - size, 28
 - square, 29
 - symmetric, 29
 - triangular, 29
 - zero, 29
- maximum
 - global, 25
 - local, 25
- mean, 34
- measurement, 49
 - scale, 49
 - unit, 50
- memory, 45
- minimum
 - global, 25
 - local, 25
- mixture, 37
- model
 - regression, 61
- module, 9
- moment, 34
 - central, 34
- multiplier
 - Floquet, 46
 - Lagrange, 25
- norm
 - of vector, 33
- number
 - complex, 19
 - rational, 19
 - real, 19
- operator, 18
 - bijjective, 18
 - Boolean, 17
 - injective, 18
 - surjective, 18
- Order, 22
- order, 22
- orthogonal, 33
- orthonormal, 33
- output, 44
- over, 18
- parabola
 - tangent, 24
- parameter, 19
 - unidentifiable, 63
- part
 - imaginary, 19
 - real, 19
- partition, 17
- Pascal, 8
- plane
 - tangent, 25
- point
 - control, 54
 - critical, 25
 - floating, 51
 - inflection, 23
 - saddle, 25
- precision, 51
- probability
 - conditional, 37
- process
 - birth and death, 42
 - branching, 42
 - continuous time, 41
 - counting, 43
 - deterministic, 41
 - discrete time, 41
 - ergodic, 43
 - interval, 43
 - Markov, 42
 - orderly, 43
 - point, 43
 - renewal, 43
 - stochastic, 41
 - Yule, 43
- product
 - direct, 28
 - inner, 33
 - matrix, 28
 - outer, 33
- range, 19
- rank, 30
 - full, 30
- rate
 - hazard, 43

- root, 27
- rule
 - chain, 24
 - de Morgan, 18
 - L'Hopital, 23
- sample, 33
- scalar, 19
- scale
 - measurement, 49
 - organization, 8
- secant, 20
- seed, 64
- sequence, 22
 - Fibonacci, 22
- series, 22
 - binomial, 22
 - geometric, 22
 - harmonic, 22
 - Maclaurin, 23
 - Taylor, 23
- set, 17
- simplex, 40
- sine, 20
- sphere, 28
- stability
 - global, 46
 - local, 46
 - neutral, 46
- state
 - steady, 41
 - transient, 41
- subset, 17
- support, 11
- system
 - non-homogeneous linear, 31
- tangent, 20
- term, 22
- test
 - power, 55
- testability, 10
- theorem
 - algebra, 27
 - Buckingham, 7
 - central limit, 36, 43
 - Leibniz, 27
 - Neyman-Pearson, 55
 - Pythagoras, 33
- theory
 - auxiliary, 1
- trace, 29
- transformation, *see* operator, 49
- transposition, 29
- union, 17
- usefulness, 3
- value
 - absolute, 19
 - average, 26
 - boundary, 44
 - eigen, 31
 - initial, 44
 - maximum, 25
 - minimum, 25
- variable
 - continuous random, 33
 - dependent, 19
 - discrete random, 33
 - extensive, 7
 - independent, 19
 - intensive, 7
 - intermediate, 24
 - random, 33
 - state, 43
- variance, 34
- vector, 29
 - eigen, 32
 - length, 33
- verification, 3