



HAL
open science

Using Automatic Differentiation to study the sensitivity of a crop model

Claire Lauvernet, Laurent Hascoet, François-Xavier Le Dimet, Frédéric Baret

► **To cite this version:**

Claire Lauvernet, Laurent Hascoet, François-Xavier Le Dimet, Frédéric Baret. Using Automatic Differentiation to study the sensitivity of a crop model. *Lecture Notes in Computational Science and Engineering*, 2012, 87, p. 59 - p. 69. 10.1007/978-3-642-30023-3_6 . hal-00766141

HAL Id: hal-00766141

<https://hal.science/hal-00766141v1>

Submitted on 17 Dec 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using Automatic Differentiation to study the sensitivity of a crop model

Lauvernet, C., Hascoët, L., Le Dimet, F.-X., and Baret, F.

Abstract Automatic Differentiation (AD) is often applied to codes that solve partial differential equations, e.g. in geophysical sciences or Computational Fluid Dynamics. In agronomy, the differentiation of crop models has never been performed since these models are more empirical than fully mechanistic, derived from equations. This study shows the feasibility of constructing the adjoint model of a crop model referent in the agronomic community (STICS) with the TAPENADE tool, and the use of this accurate adjoint to perform some sensitivity analysis. This paper reports on the experience from AD users of the environmental domain, in which AD usage is not very widespread.

Key words: adjoint mode, agronomic crop model, sensitivity analysis

1 The application domain: the agronomic crop model STICS

STICS [2, 3] is a crop model with a daily time step. Its main aim is to simulate the effects of the physical medium and crop management schedule variations on crop production and environment at the field scale. From the characterization of climate, soil, species and crop management, it computes output variables related to yield in

Claire Lauvernet

Irstea, UR MALY, 3 bis quai Chauveau - CP 220, F-69336 Lyon, France (previously at INRA Avignon, France) claire.lauvernet@irstea.fr

Laurent Hascoët

INRIA, Sophia-Antipolis, France Laurent.Hascoet@inria.fr

François-Xavier Le Dimet

Université de Grenoble, France Francois-Xavier.Le-Dimet@imag.fr


Frédéric Baret

INRA, Avignon, France baret@avignon.inra.fr

terms of quantity and quality, environment in terms of drainage and nitrate leaching, and to soil characteristics evolution under cropping system¹.

The two key output variables simulated by STICS that we will need in this paper are the Leaf Area Index (LAI) and the biomass. The LAI is the total one-sided area of leaf tissue per area of ground surface (unitless). This is a canopy parameter that directly quantifies green vegetation biomass. As the leaves are considered to be the main interfaces with the atmosphere for the transfer of mass and energy [16], the LAI indirectly describes properties such as potential of photosynthesis available for primary production, plant respiration, evapotranspiration and carbon flux between the biosphere and the atmosphere, and gives evidence of severely affected areas (fires, parasites. . .). Because it is the most observable canopy parameter by remote sensing, the LAI is very commonly used e.g., in crop performance prediction [7], in models of soil-vegetation-atmosphere [15], in crop models [2, 3], in radiative transfer models [20]. Its values can range from 0 for bare soil to 6-7 for a crop during its life cycle, and up to 15 in extreme cases (tropical forests).

Fig. 1 Simplistic scheme of the stages simulated by the STICS model on dynamics of LAI.



STICS simulates the crop growth from sowing to harvest, focusing on the evolution of the LAI at a few selected [2] vegetative stages shown on Fig. 1. These stages involve process thresholds, accounting for some of the differentiation problems described in Sect. 3.2. For a wheat crop, the main phenological stages are known as *ear at 1 cm*, *heading*, *flowering*, and *maturity*. In this work we do not simulate grain yield but only the total biomass. As we focus on the LAI, we only consider the vegetative stages namely: LEV (emergence or budding), AMF (maximum acceleration of leaf area index, equivalent to *ear at 1 cm*), LAX (maximum LAI i.e. end of leaf growth), and SEN (start of net senescence).

2 Sensitivity Analysis

A model is a more or less realistic or biased simplification of the state variables it simulates. This is especially true for agronomic models, since the functioning of vegetation is not *a priori* described by exact equations: agronomic models attempt to predict the behavior of the crop by incremental improvements of the simulation code, based on observations made on the field and then published by specialists. Thus, in some parts of the model, this empirical approach is not based on the equations of some underlying physics or chemistry. Sensitivity analysis, which studies

¹ http://www.avignon.inra.fr/agroclim_stics_eng

the impact of perturbing the control parameters on the model output, gives insights useful to improve or even simplify the model. Sensitivity analysis requires two essential ingredients:

- A model: $F(X, K) = 0$, where X is the state variable (LAI, biomass . . .) and K the control variables (parameters, forcing variables . . .). F is a differential operator *a priori* non-linear finite-dimensional, that describes implicitly the evolution of X for a given K . We assume that the system has a unique solution $X(K)$. In this study, what we call the model is exactly the STICS computer program.
- A response function G which combines one or more elements of X into a scalar value, e.g. the final value or the integral over time of an output.

The problem is to evaluate the sensitivity of G with respect to K or in other words the gradient of G with respect to K . With the help of the adjoint model, computing the gradient takes only 2 steps: run the direct model once for the given K , then solve the adjoint model once [12]. The classical justification is:

$$\nabla G = \frac{dG}{dK} = \left(\frac{dG}{dX} \cdot \frac{dX}{dK} \right)^t = \left(\frac{dX}{dK} \right)^t \cdot \left(\frac{dG}{dX} \right)^t$$

where we observe that $\frac{dG}{dX}$ is easily computed from the definition of G alone and the product of $\frac{dX}{dK}^t$ with a vector is achieved by feeding this vector to the adjoint code of STICS, produced by the adjoint mode of Automatic Differentiation.

Sensitivity analysis using an adjoint model is the only way to calculate formally the gradient of the response function at a cost that does not depend on the size of K . It is particularly suitable when the number of entries K is large compared to the size of the response function G [14, 13].

One can also compute the gradient accurately with tangent-linear differentiation, at a cost that is proportional to the size of K . The other sensitivity methods only approximate the gradient: finite difference approximation of the gradient require extensive direct model computations [4]. Stochastic sampling techniques require less mathematical insight as they consist (roughly speaking) in exploring the space of control to determine an overall global sensitivity [18, 10]. Their cost grows rapidly with the dimension of K . These methods have been widely applied to the agronomic models and in particular on STICS [9, 17, 19].

If in many cases, the response function G is a differentiable function of K , it can happen that the model is driven by thresholds e.g., the code uses a lot of branches. Theoretically, a piecewise continuous function is not continuously differentiable, but it has right- and left-derivatives. Differentiation of such a code can only return a sub-gradient. Actually, the methods that do not rely on derivatives (divided differences, stochastic, . . .) behave better in these cases, although they remain expensive. In practice, this problem is not considered serious as long as the local sensitivity is valid in a neighborhood of the current K .

3 Automatic Differentiation of STICS

3.1 *The TAPENADE Automatic Differentiation tool*

TAPENADE [8] is an Automatic Differentiation (AD) tool based on source transformation. Given a source program written in FORTRAN, TAPENADE builds a new source program that computes some of its derivatives. In “tangent” mode, TAPENADE builds the program that computes directional derivatives. In “adjoint” mode, TAPENADE builds the program that computes the gradient of the output with respect to all input parameters.

Considering the complete set of derivatives of each output with respect to each input, i.e. the Jacobian matrix of the program’s function, the tangent mode yields a column of the Jacobian whereas the adjoint mode yields a row of the Jacobian. Therefore in our particular case where the output is a scalar G , one run of the adjoint code will return the complete gradient. In contrast, it takes one run of the tangent mode per input to obtain the same gradient. Although we will experiment with the two modes, the adjoint mode fits our needs better.

However, the adjoint mode evaluates the derivatives in the inverse of the original program’s execution order. This is a major difficulty for large programs such as STICS. The AD model copes with this difficulty by a combination of storage of intermediate values and duplicated evaluation of the original program, at a cost in memory and execution time. In TAPENADE, the strategy is mostly based on storage of intermediate values, combined with the storage/recompute tradeoff known as checkpointing, applied automatically at each procedure call.

3.2 *STICS adjoint : the pains and sufferings of an AD end-user*

The STICS model being written in FORTRAN 77, TAPENADE can in theory build its adjoint. However, there were shortcomings with the early versions of TAPENADE, before 2005. Later versions brought notable improvements but we believe it is worth describing the main problems that we encountered at these early stages.

AD allows for instructions which the symbolic differentiation systems cannot process. It also provides a real gain in computational time. However, a few good programming practices are recommended: the input parameters involved in derivatives must be clearly identified and if possible separate from the other variables. The same holds for the outputs to be differentiated. The precision level of all floating point variables must be coherent, especially for validation purposes: if the chain of computation is not completely “double precision”, then the divided difference that is used to validate the analytic derivatives will have a poor accuracy, validation will be dubious and may even fail to detect small errors in the differentiated code.

Validation helped us detect small portability problems in STICS. As divided differences requires to call STICS twice, we discovered that two successive calls to

STICS apparently with the same inputs gave different results. In fact the first call was different from all the others, which pointed us to a problem of hidden uninitialized remanent global. Fixing this gave us correct divided differences, and a more portable STICS code.

More specifically to this agronomy application, we had problems with the high number of tests and other conditional jumps in an average run. In more classical situations of Scientific Computing, programs are derived from mathematical equations, typically a set of ODE's or PDE's. This forces some regularity into the code that discretizes and solves these equations: even if branches do occur, they rarely introduce discontinuity and the derivative itself often remains continuous. In our application, the program itself basically is the equation. The model evolves by introducing by hand new subcases and subdivisions, i.e. more tests. If this evolution is not made with differentiation in mind, it may introduce sharp discontinuities that do not harm the original code but make it non-differentiable. It took time to replace faulty branches with a cleaner, differentiable implementation. On the other hand, users agreed that this resulted in a better code.

Still, the number of branches in the STICS model is very large: thresholds, conditions, loops, and other control, all are tests that the adjoint code must remember to run backwards. STICS consumes an unusually large memory for that. Until recently, TAPENADE did not store this control efficiently, using a full INTEGER value to store only a boolean in general.

Checkpointing the time stepping was difficult. Before binomial checkpointing [5] was implemented in TAPENADE, we had to split the main time loop of 400 iterations into two nested loops of 20 iterations each, and place these two loops into two new subroutines to force checkpointing. This tedious manipulations are now spared with the new TAPENADE directives for binomial checkpointing.

More than five years after this sensitivity study, both STICS and TAPENADE have evolved. The latest version 6 of STICS is more readily differentiable than before. TAPENADE 3.6 had several bugs fixed and more importantly provides a set of user directives to control checkpointing better. These checkpointing directives are also the answer to the serious performance problem discussed in Sect. 3.3.

3.3 Validation of the adjoint model

Validation was performed in two steps as usual, and for several directions of perturbation. First, the tangent derivatives were compared with divided differences, and they agreed up to the 8th decimal for an increment of 10^{-8} in the one-sided divided difference. Second, the adjoint derivatives were compared with the tangent derivatives ("dot-product" test [6]) and they agreed up to the 14th decimal. At the time of the study, the run times were:

Direct model : 0.21s Tangent model : 0.39s Adjoint model : 30.96s

The run time of the adjoint code is much higher than the customary 5- to 10-fold slowdown. The problem was left to the TAPENADE developers to go on with the sensitivity study. Identifying its causes was hard, and pointed to the need for specific profiling tools for adjoint codes. Profiling instructions must be inserted by the AD tool itself, and tools are missing to help interpret the profiling results. Eventually, the problem was found to come from the systematic checkpointing on procedure calls on a chain of 4 nested procedure calls, each of them doing little else than calling the next nested call, cf Fig. 2. Checkpointing [6] one call to \mathbb{P} reduces the peak memory used by the adjoint. This reduction is roughly proportional to the run-time of \mathbb{P} . On the other hand, it costs one extra run of \mathbb{P} , plus some memory (a “snapshot”) to restore the execution state. Checkpointing nested calls causes an increasing number of extra runs. This is inherent to the approach and beneficial in




Fig. 2 The cost of checkpointing long chains of nested calls

general, but is a waste for procedures that are nearly empty shells around a deeper call. In our case, the problem was amplified by the size of a very big work array that was restored at each checkpoint. The answer is to deactivate checkpointing on the calls to the “empty shell” procedures. This is known as the “split” mode of adjoint AD [6], and is sketched on the right of Fig. 2. This required development in TAPENADE, plus new directives (`$AD NOCHECKPOINT`) to let the user trigger this split mode on selected procedure calls. Conversely in other cases, it is useful to trigger checkpointing on pieces of a procedure, and TAPENADE new directives (`$AD NOCHECKPOINT`) and (`$AD NOCHECKPOINT`) let the user do that. This results in the following times obtained with TAPENADE 3.6:

Direct model : 0.22s Tangent model : 0.52s Adjoint model : 0.86s

4 Results : sensitivity analysis of STICS

We decided to compute the gradients of two response functions G : LAI and biomass, and more precisely their integrals over the simulation time from sowing to harvest. These response functions capture well the growth dynamics.

$$G_{LAI} = \sum_{i=1}^T LAI(t_i) \quad G_{biomass} = \sum_{i=1}^T biomass(t_i)$$

4.1 Selection of input parameters for sensitivity analysis of output variables

For this feasibility study, the control variables correspond to wheat crops from the Danube's plain in Romania in 2000-2001² [1]. The gradient was calculated with respect to the following input parameters³: for LAI, we chose the varietal parameters acting on the dynamics of LAI, and *dlaimaxbrut* that strongly characterizes the aerial growth. Parameters were adapted to the ADAM database, including the variety of wheat (Flamura) used here for its particular cold resistance. For biomass, efficiencies at three important phases of the cycle of wheat (juvenile phases, vegetative and grain filling) and *vmax2* were chosen following the experience accumulated by users of the crop model. Table 1 describes the role of these parameters, and their values for this sensitivity study.

Table 1 Parameter role and values for the ADAM conditions

parameter	definition	value
<i>dlaimaxbrut</i>	maximum rate of gross leaf surface area production	0.00044
<i>stlevamf</i>	cumulated development units between the LEV and AMF stages	208.298
<i>stamflax</i>	cumulated development units between AMF and LAX stages	181.688
<i>jvc</i>	days of vernalisation (cold days needed to lift)	35
<i>durvieF</i>	lifespan of a cm of adult leaf	160
<i>adens</i>	compensation between number of stems and plants density	-0.6
<i>efcroijuv</i>	maximum growth efficiency during juvenile phase (LEV-AMF)	2.2
<i>efcroiveg</i>	maximum growth efficiency during vegetative phase (AMF-DRP)	2.2
<i>efcroirepro</i>	maximum growth efficiency during grain filling phase (DRP-MAT)	4.25
<i>vmax2</i>	maximum rate of nitrate absorption by the roots	0.05

4.2 Sensitivity results of LAI and biomass

One goal of this sensitivity study was to establish the hierarchy of influent parameters. Therefore Fig. 3 shows the 10 influences normalized as percentages, totalling 100%. Among the 10 selected, the most influential parameters on the LAI are *adens*

² ADAM experiment (Data Assimilation through Agro-Modelling). Project and database at <http://kalideos.cnes.fr/spip.php?article68>

³ All the parameters of STICS are described in http://www.avignon.inra.fr/agroclim_stics_eng/notices_d_utilisation

(47%), *dlaimaxbrut* (21%), *stlevamf* (17%), *jvc* (10%), and finally *stamflax* (2%). *adens* represents the ability of a plant to withstand increasing densities, and since it depends on the species and varieties, its influence may be particularly strong for this type of wheat and less for other crops. For biomass, we observe that the hierarchy is modified by the strong influence of the efficiency *efcroiveg* (maximum growth efficiency during vegetative phase) which is similar to that of *adens* (27%). This means that we can ignore the estimate of *efcroiveg* if we only want to assimilate LAI data, but absolutely not if we need to simulate biomass. *stlevamf* and *dlaimaxbrut* are of similar importance (14 and 12%). Finally, there is a relatively low sensitivity (5% and 3%) of biomass integrated over the life cycle to the other two parameters of efficiency *efcroirepro* and *efcroijuv*, meaning that the biomass is not so dependant on the juvenile and the grain filling phases but essentially on the vegetative phase. The fact that only the integral over the entire cycle was studied involves a very small influence of the parameters *efcroirepro* and *efcroijuv*, as opposed to *efcroiveg*. These efficiencies with a small influence matter only during short phenological stages: only a sensitivity study restricted to these stages can modify the hierarchy of influent parameters, opening the way to estimation of these low-influence parameters [17]. LAI is actually dependant on 4 parameters and biomass on 5 on the 10 tested, which will help the user concentrate on these and estimate them better. Uncertainty on the other parameters is of relatively smaller importance.

5 Conclusion and outlook

This case study illustrates the interest of AD for sensitivity analysis of agronomic models. Coupled with other models, for example radiative transfer model [11], it will allow to assimilate remote sensing data into crop models by using the adjoint to minimize the discrepancy cost function. This work shows the feasibility of applying and developing variational methods in agronomy, in the same way as in oceanography or meteorology.

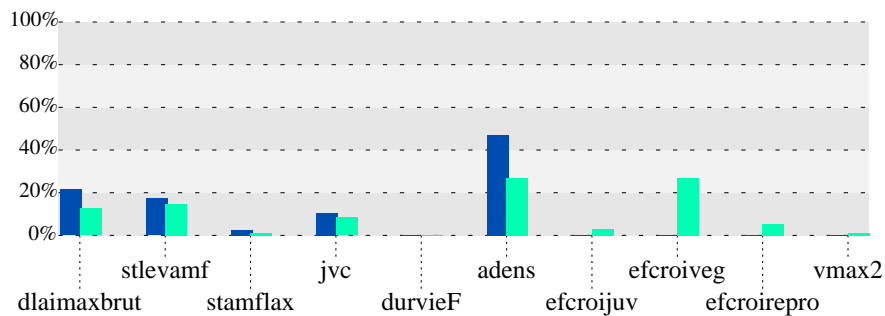


Fig. 3 Relative sensitivity (%) to selected STICS parameters of output variables LAI (left) and biomass (right) computed by the adjoint.

For the agronomic community, the adjoint model of STICS is an interesting tool to perform sensitivity analysis since it requires the calculation only once for each agro-pedo-climatic situation. The most difficult work is the differentiation of the model, which must be done only once, and with the help of AD tools that keep improving. However, the local sensitivity analysis is valid only in a small neighborhood and the hierarchy of sensitivities may vary under different conditions. These results are only a first step. Following work could concentrate on:

1. a “multi-local” sensitivity analysis, keeping the crop management and climate of the base ADAM, but letting the parameters vary in a given range. This would require many runs of the adjoint modes on a representative sample of possible parameter values. This would return a parameter hierarchy with a more general validity.
2. an application of this analysis to other conditions (climate, soil...) to see whether the hierarchy is preserved in general. Extending to other varieties is also important. Actually, it seems unlikely that this hierarchy is preserved since the change of climate and soil conditions may rapidly hit limiting factors (stress for the plant) and thus modify the parameters influence.
3. a study of the sensitivity at selected phenological stages of the cycle to study the effect of variables temporally valid (especially efficiency) on the general hierarchy.

The adjoint code is able to compute the sensitivities of one response function to all parameters in just one run. There are more parameters to STICS than the 10 we have selected for this sensitivity study. Looking at the influence of all parameters will guide the attention of STICS users on some parameters and modules, according to the users' objectives. Sensitivity study is a preliminary to parameter estimation: many of these agronomic parameters (yield, balance . . .) are not directly observable by remote sensing. On the other hand the outputs (biomass) can be measured. The adjoint of the model, by returning the gradient of any discrepancy cost function, is the key to estimate these hidden agronomic parameters from the ones we can measure.

Acknowledgements This study was conducted thanks to a grant provided by CNES within the ADAM project (<http://kalideos.cnes.fr/spip.php?article68>), during the Phd of the first author at INRA Avignon and the University of Grenoble.

References

1. Baret, F., Vintila, R., Lazar, C., Rochdi, N., Prévot, L., Favard, J., de Boissezon, H., Lauvernet, C., Petcu, E., Petcu, G., Voicu, P., Denux, J., Poenaru, V., Marloie, O., Simota, C., Radnea, C., Turnea, D., Cabot, F., Henry, P.: The adam database and its potential to investigate high temporal sampling acquisition at high spatial resolution for the monitoring of agricultural crops. *Romanian Agricultural Research* **16**, 69–80 (2001)
2. Brisson, N., Mary, B., Ripoche, D., Jeuffroy, M.H., Ruget, F., Nicoullaud, B., Gate, P., Devienne-Barret, F., Antonioletti, R., Durr, C., Richard, G., Beaudoin, N., Recous, S., Tayot,

- X., Plenet, D., Cellier, P., Machet, J.M., Meynard, J.M., Delecolle, R.: STICS : a generic model for the simulation of crops and their water and nitrogen balances. I: theory and parameterization applied to wheat and corn. *Agronomie* **18**(5-6), 311–346 (1998)
3. Brisson, N., Ruget, F., Gate, P., Lorgeou, J., Nicoullaud, B., Tayot, X., Plenet, D., Jeuffroy, M.H., Bouthier, A., Ripoche, D., Mary, B., Justes, E.: STICS: a generic model for simulating crops and their water and nitrogen balances. II: model validation for wheat and maize. *Agronomie* **22**(1), 69–92 (2002)
 4. Castaing, W., Dartus, D., Le Dimet, F.X., Saulnier, G.M.: Sensitivity analysis and parameter estimation for distributed hydrological modeling: potential of variational methods. *Hydrol. Earth Syst. Sci.* **13**(4), 503–517 (2009)
 5. Griewank, A.: Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. *Optimization Methods and Software* **1**, 35–54 (1992)
 6. Griewank, A., Walther, A.: *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, 2nd edn. No. 105 in *Other Titles in Applied Mathematics*. SIAM, Philadelphia, PA (2008). URL <http://www.ec-securehost.com/SIAM/OT105.html>
 7. Guérif, M., Houlès, V., Makowski, D., Lauvernet, C.: Data assimilation and parameter estimation for precision agriculture using the crop model STICS. In: D. Wallach, D. Makowski, J.W. Jones (eds.) *Working with dynamic crop models: evaluating, analyzing, parameterizing and using them*, chap. 17, pp. 391–398. Elsevier (2006)
 8. Hascoët, L., Pascual, V.: TAPENADE 2.1 user's guide. Rapport technique 300, INRIA, Sophia Antipolis (2004). URL <http://www.inria.fr/rrrt/rt-0300.html>
 9. Houlès, V., Mary, B., Guérif, M., Makowski, D., Justes, E.: Evaluation of the ability of the crop model stics to recommend nitrogen fertilisation rates according to agro-environmental criteria. *Agronomie* **24**(6), 339–349 (2004)
 10. Ionescu-Bujor, M., Cacuci, D.G.: A comparative review of sensitivity and uncertainty analysis of large-scale systems. I: deterministic methods. *Nuclear science and engineering* **147**(3), 189–203 (2004)
 11. Lauvernet, C., Baret, F., Hascoët, L., Buis, S., Le Dimet, F.X.: Multitemporal-patch ensemble inversion of coupled surface-atmosphere radiative transfer models for land surface characterization. *Remote Sens. Environ.* **112**(3), 851–861 (2008)
 12. Le Dimet, F.X., Ngodock, H.E., Navon, I.M.: Sensitivity analysis in variational data assimilation. *J. Meteorol. Soc. Japan* pp. 145–155 (1997)
 13. Le Dimet, F.X., Talagrand, O.: Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A* **38A**(2), 97–110 (1986)
 14. Lions, J.L.: *Optimal control of systems governed by partial differential equations*. Springer-Verlag (1968)
 15. Olioso, A., Inoue, Y., Ortega-FARIAS, S., Demarty, J., Wigneron, J., Braud, I., Jacob, F., Lecharpentier, P., Ottl, C., Calvet, J., Brisson, N.: Future directions for advanced evapotranspiration modeling: Assimilation of remote sensing data into crop simulation models and SVAT models. *Irrigation and Drainage Systems* **19**(3-4), 377–412 (2005)
 16. Rosenberg, N.J., Blad, B.L., Verma, S.B.: *Microclimate: the biological environment*. Wiley-Interscience (1983)
 17. Ruget, F., Brisson, N., Delecolle, R., Faivre, R.: Sensitivity analysis of a crop simulation model, STICS, in order to choose the main parameters to be estimated. *Agronomie* **v. 22**(2) p. **133-158** (2002)
 18. Saltelli, A., Chan, K., Scott, E.M.: *Sensitivity Analysis*. Wiley (2000)
 19. Varella, H., Guérif, M., Buis, S.: Global sensitivity analysis measures the quality of parameter estimation: The case of soil parameters and a crop model. *Environmental Modelling and Software* **25**(3), 310 – 319 (2010)
 20. Verhoef, W.: Light scattering by leaf layers with application to canopy reflectance modeling: The SAIL model. *Remote Sensing of Environment* **16**(2), 125–141 (1984)