

Adaptive hybrid physical/neural network modeling, and its application to greenhouse climate optimization.

Raphael Linker*
Dept. of Agricultural Eng.
Technion
Haifa, 32000 Israel

Ido Seginer
Dept. of Agricultural Eng.
Technion
Haifa, 32000 Israel

Per-Olof Gutman
Dept. of Agricultural Eng.
Technion
Haifa, 32000 Israel

Abstract

This paper focuses on adaptive modeling of non-linear systems which operate in slowly changing environments. Due to an inability to control the environment, a large amount of data spanning the whole feasible input space can not be collected over a reasonably short period of time. As a result, modeling such systems with neural networks, which usually have poor extrapolation properties, might lead to poor results. In order to avoid these poor extrapolations, hybrid physical/neural network models are used. Such models are formed by the combination, in parallel, of a physical model (approximately valid over the whole space), and a radial basis function (RBF) network, which provides localized predictions only where training data are available.

In this paper, the hybrid modeling approach is extended by adapting the RBF network on-line, so that the region over which the network is valid grows over time. In order to simplify the RBF training, the centers are located on a fixed rectangular grid: if necessary, a new center is added at the grid point closest to the new datapoint. This approach also allows for keeping the RBF width constant and equal for all the centers.

The problem of discarding some of the data, so that the database does not become prohibitively large, is also addressed. In order to avoid 'forgetting' previously modeled regions of the space, datapoints are not discarded based solely on their 'age'. Rather, the number of datapoints near each center is limited, and when this limit is reached, the oldest datapoint associated with that center is discarded.

Greenhouse climate modeling, and its use for climate optimization, is presented as an illustration of the method.

1 Introduction

Greenhouses provide a protected environment in which crops can be grown under a tightly controlled climate. Since the control of the greenhouse climate is costly, its optimization has been studied by several authors (Seginer *et al.*, 1986; Gutman *et al.*, 1993; Ioslovich *et al.*, 1995; Chalabi and Zhou, 1997). Optimization requires adequate models of the greenhouse and of the

*Email: linkerr@tx.technion.ac.il

crop. The present work focuses on the greenhouse model, while the crop model is taken from the literature.

A situation in which the greenhouse manufacturer provides not only the greenhouse but also a model of that greenhouse (such as the model used to design the greenhouse), is considered. Such a model does not include location-dependent effects such as orientation, surrounding buildings, local wind, etc.. The goal of the present work is to answer the following questions: How can the original (manufacturer) model be adapted to the specific location? How much time is required for this adaptation, and by how much does it improve the greenhouse operation?

The greenhouse models described in the literature can be classified into two categories: physical models, and black-box models. Physical models (Takakura *et al.*, 1971; Kindelan, 1980; Avissar and Mahrer, 1982; Boulard and Baille, 1993) are based on energy balance equations between the indoor air, the outdoor air, and the greenhouse soil (in which heat storage takes place). In order to be accurate, such models require the calibration of a relatively large number of parameters such as cover transmissivity and heat transfer coefficient, air-soil heat transfer coefficient, soil heat capacity, etc.. The calibration of such parameters requires time-consuming dedicated experiments. In order to avoid the calibration of these 'physical' parameters, Seginer *et al.* (1994) suggested the use of black-box neural network models. The use of such models was reported also by Kok *et al.* (1994), and Linker *et al.* (1998b) extended their use to the determination of optimal temperature and CO₂ concentration setpoints. While these studies showed that neural networks could be trained to model greenhouse climate accurately, Linker *et al.* (1998b) emphasized the problems associated with the poor extrapolation properties of such models. Due to the strong influence of the outdoor weather on the greenhouse climate, the input vector of a greenhouse model includes uncontrollable variables, which makes the collection of training data covering the whole operating domain virtually impossible over a reasonable period of time. Furthermore, the high dimensionality of the input space makes it difficult to detect input vectors which are outside the training domain and cause the neural network to extrapolate.

Hybrid physical/neural network modeling was introduced by Thompson and Kramer (1994) as a mean to overcome, at least partially, the poor extrapolation properties generally associated with neural networks. Such a hybrid modeling can be interpreted as a way to include prior knowledge in a neural network (NN). Basically, NN models are black-box models, and as such, do not contain prior knowledge about the system to be modeled, except for the knowledge contained in the choice of the input and output variables (Rudolph, 1997). The absence of prior knowledge explains, at least partially, the poor extrapolation properties of the NN models. In models based on physical understanding of the system (physical models), the prior knowledge is automatically reflected by the structure of the model. This structure acts as a constraint on the mapping that the model is able to create between the input and output variables. In such cases, the model structure (and its complexity) directly affects the ability of the model to approximate correctly the real system.

For neural networks, this constraint does not exist, and the NN is free to create any input-output mapping, as long as this mapping provides good predictions for the training/test data. An indirect way to constraint the input-output mapping (or to introduce prior knowledge in the NN) is to create additional training points using a physical model. This augmented-data-set approach was introduced by Tsen *et al.* (1996), and an application of this method was presented by Milanic *et al.* (1997). However, this method has the following drawbacks: (1) the data set may become very large, which slows the training process, and (2) if few experimental points are available, these may be 'lost' among the artificial points, and the NN will replicate the physical model.

Several authors have tried to restrict the mapping created by neural networks by imposing constraints either on the model structure (design approaches) or on the values that the parameters (weights) are allowed to take (training approaches) (Joerding and Meador, 1991; Thompson and Kramer, 1994). The latter can enforce monotonicity, convexity (or concavity), or smoothness of the input-output mapping (e.g. Bishop (1993)). However, such methods are mathematically and computationally complex, and their use remains very limited. Design approaches use prior knowledge as the basis for the selection of the type of network, activation functions, inclusion of non-network parametric models, or modularization of the network architecture. Modular design approaches result in partially connected networks. Each subnetwork corresponds to a subprocess of the process modeled. Methods for the inclusion of non-network parametric models, which lead to hybrid models, are reviewed below.

1.1 Series hybrid modeling

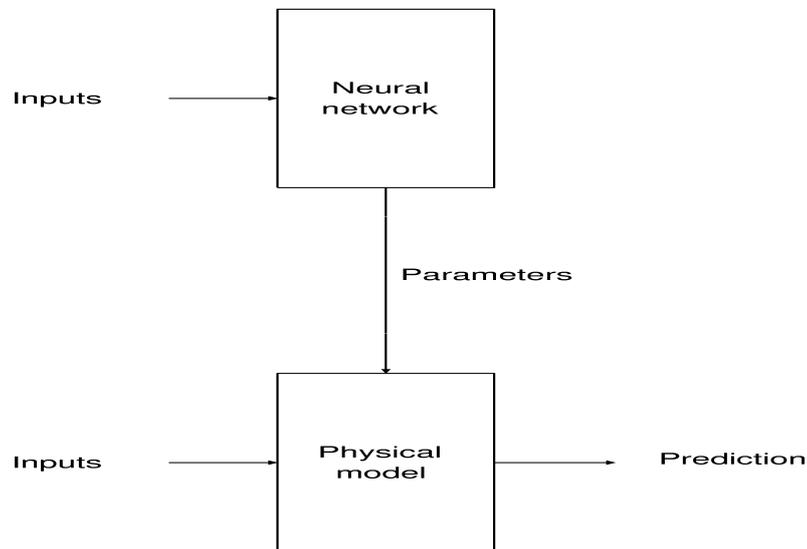


Figure 1: Series hybrid modeling

A schematic representation of the series hybrid modeling approach (Thompson and Kramer, 1994) is presented in Figure 1, where a NN is used to predict one or several parameters of a physical model. In this manner, the structure of the physical model is maintained, but the parameters are expressed as (complex) functions of the environment or working point. It must be noted that the output(s) of the NN are parameter values, which are not available to train the NN straightforwardly. However, it is possible to train the NN by back-propagating the prediction error through the physical model, using the model's Jacobian (such an approach is similar to learning with distal teacher (Jordan and Rumelhart, 1992), and inverse model neural control (Tanomaru and Omatu, 1992)). This approach was also presented by Psychogios and Ungar (1992), and applications were presented by Feyo de Azevedo *et al.* (1997), Piron *et al.* (1997), and Wilson and Zorsetto (1997). All these studies reported extrapolation performances superior to the ones of either standard sigmoid NNs or pure physical models. However, series hybrid modeling presents only a partial solution to the extrapolation problem of neural networks: NN extrapolation is not avoided, but rather the NN might have to extrapolate the parameter

values instead of the model output itself.

1.2 Parallel hybrid modeling

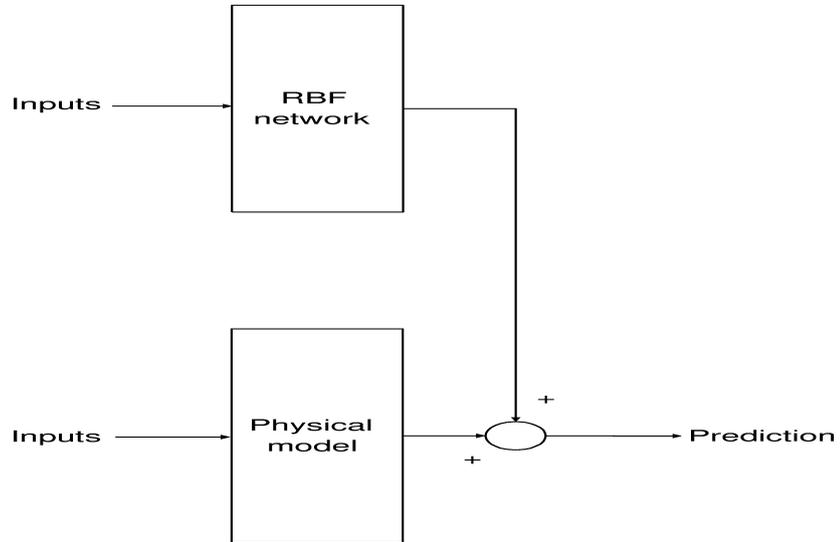


Figure 2: Parallel hybrid modeling

Parallel hybrid modeling (Thompson and Kramer, 1994) involves the arrangement of a physical model, which is approximately valid over the whole input space, and a radial basis function (RBF) neural network (Figure 2). The predicted output is the sum of the physical model and RBF outputs, so that the RBF contribution can be viewed as a correction to the physical model. If the RBF network is trained by assigning centers only where training data are available, the RBF output differs from zero only inside (and near) the training domain. In this manner, the physical model provides a first approximation, which is improved wherever training data is available.

1.3 Focus of this paper

This work extends the parallel hybrid modeling approach by including on-line adaptation of the RBF component of the model. The motivation is that, in many practical applications, an initial model of the system is available. While such a model might be rather coarse and not sufficiently accurate for optimal (or satisfactory) operation, it might be sufficient for starting the system operation. The data collected on-line can then be used to improve the model, in the present case by adapting the RBF component of the hybrid model. At this stage, the physical component of the hybrid model, which provides first approximation (default) predictions over the whole space, is not adapted.

2 Adaptive hybrid modeling

The training phase of a RBF network requires the calculation (or *a priori* determination) of the centers' position, the widths of the activation functions, and the output weights. The original

algorithm for training RBF networks was developed by Moody and Darken (1989). According to this algorithm, the centers and widths of the neurons are determined in a self-organizing manner, so that centers are placed only in regions of the input space where datapoints are present. The standard k -means clustering algorithm is used to find the location of the centers, and the neurons widths are determined using some P nearest neighbors heuristic. The output weights are determined using the supervised least-mean-square (LMS) rule.

Chen *et al.* (1992) adapted the original RBF training algorithm of Moody and Darken (1989) to non-linear system identification. Their main contribution was to implement the k -means algorithm recursively, so that, if the data distribution is changing, the centers' distribution can follow the data. However, neither Moody and Darken (1989) nor Chen *et al.* (1992) did address the problem related to the required number of neurons in a systematic manner. Platt (1991) proposed an ever-growing network in which new centers are added whenever an unusual (or new) input-output pair is presented to the network. Novelty is determined by two factors: (1) the datapoint is far from the existing centers, and (2) the difference between the desired and predicted output is large. In such cases, a new center is added at the location of the new datapoint. Whenever there is no need to allocate a new center, the LMS algorithm is used to adjust the centers' positions and widths, and the output weights, in order to reduce the prediction error. Simulations of adaptive time series prediction showed that the network resulting from Platt's method attains accuracy comparable to that of a back-propagation trained network, but requires fewer computations.

Yingwei *et al.* (1997) further extended the idea of time-varying system identification by not only adding centers, but also removing 'obsolete' ones. Centers which consistently make little contribution to the output are assumed to model dynamics which are no longer present in the system, and are removed. Simulation results showed that the networks obtained with this add/drop algorithm were far more compact than the networks obtained with Chen *et al.* (1992) algorithm.

An approach somewhat similar to that of Platt (1991) is used here to adapt the RBF component of a hybrid model. Instead of adding the centers at the location of the new datapoints, the centers are added on a rectangular grid: A new center is added at the grid point closet to the new datapoint. This is illustrated on a two-dimensional example in Figure 3. In this figure, the red '+' denote old datapoints, with which are associated centers indicated by green '□'. When the new datapoints denoted by cyan '*' become available, new centers, denoted by blue 'o', are added. This grid approach allows for keeping the width of the Gaussian activation function constant, and equal for all the centers, which simplifies the computations. Also, novelty is determined only by the distance of the datapoint from the existing centers, regardless of the prediction error. In the second phase of the adaptation procedure, prediction errors are minimized by computing the linear output weights using the least squares algorithm. Note that since a zero contribution of the RBF network is desired outside of the training domain, the bias of the output neuron is not adapted (it is determined by the output scaling, if any).

3 Database formation

The previous section was concerned with the adaptation of the model parameters. However, the problem regarding which data should be used to train (or adapt) the model was not considered. This is a classic problem in adaptive system modeling. As time progresses, the amount of data available becomes prohibitively large, and some of the data have to be discarded. In general, the oldest data are discarded first, which implicitly assumes that the goal of the adaptation is to obtain a good model around the current working point. Such a selection of the training data

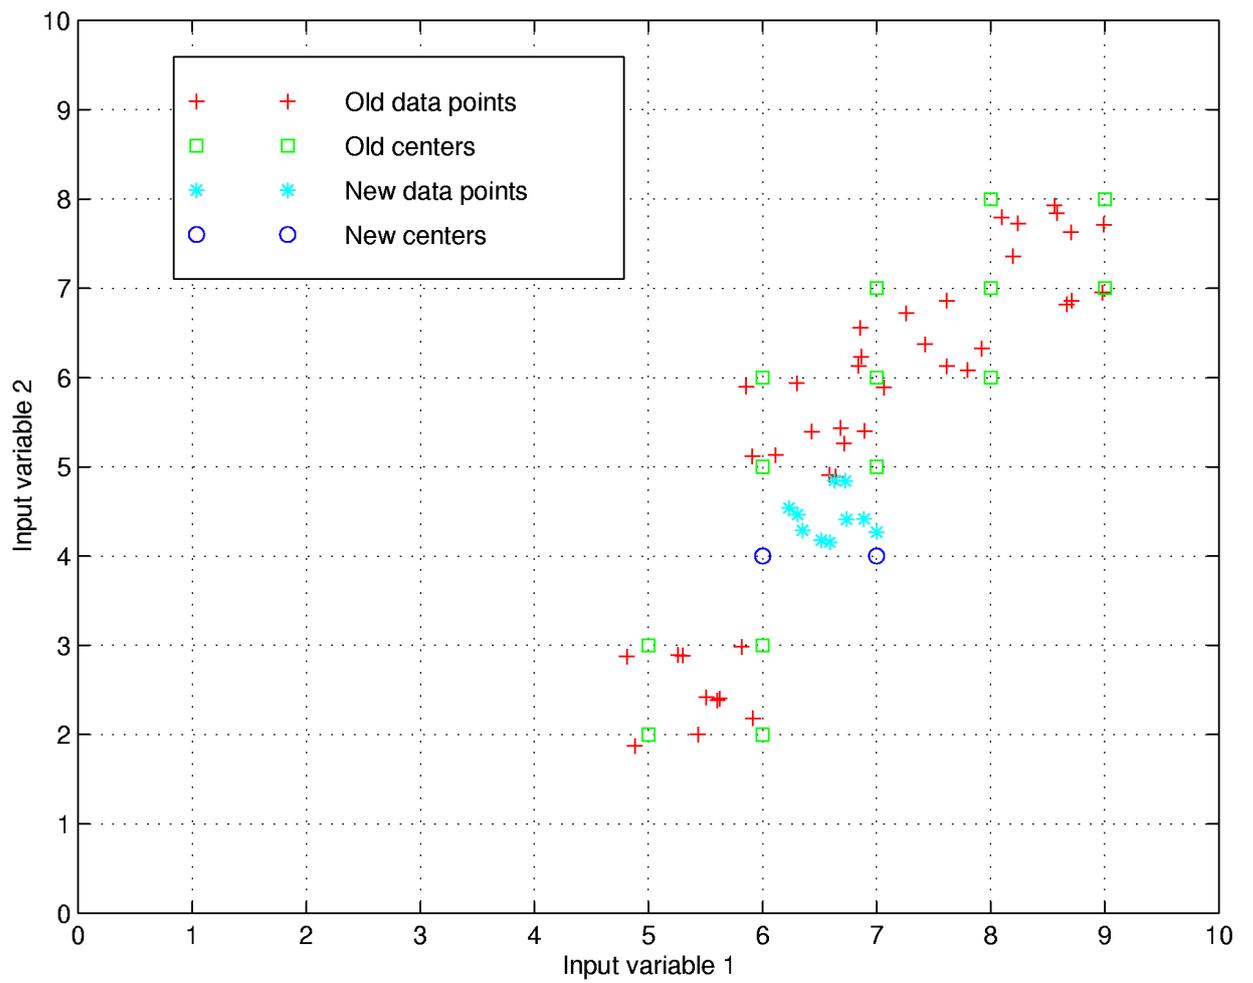


Figure 3: Grid positioning of the radial basis function centers for a two-dimensional example.

is not suitable for the present application, which is concerned with enlarging the domain over which the RBF network is valid.

Due to the locality of the Gaussian activation function, each datapoint is associated primarily with one center (the closest one). Such a center also corresponds to a specific region of the input space, or in other words, to a given working point. Therefore, the database is formed according to the following rule: A limit is imposed on the number of datapoints associated with each center. When this limit is attained, the oldest datapoint associated with this center is discarded. This can be viewed as a local implementation of the standard method described previously. For a given center, or a given working point, only the newest data are used. However, the addition of new data resulting from a change of working point does not cause the model to ‘forget’ the information recorded around old working points. This approach is somewhat opposite to that of Yingwei *et al.* (1997), because the problems considered are fundamentally different. Yingwei *et al.* (1997) was concerned with modeling of dynamic systems, and old (obsolete) centers were assumed to correspond to dynamics no longer present in the system. In the present case, not the system itself, but rather the environment in which it operates, is changing.

4 Application to greenhouse climate modeling and optimization

Greenhouse climate modeling is used as an illustration of the method. Due to the lack of experimental data, a simulation study, in which the ‘real’ greenhouse is described by the models detailed in Section 4.1, is used.

The optimal control for the daytime (solar radiation higher than 10 [W/m²]) weather recorded at the Technion, Haifa, Israel, between April 25 and October 17, 1997 is determined using the hybrid models described in Section 4.3. The optimization is performed every half an hour, using averaged weather measurements. The instantaneous criterion to be maximized is the difference between the added dry matter and the CO₂ enrichment cost (note that the crop model is assumed to be known):

$$j = K_g G - K_c R. \quad (1)$$

The crop photosynthesis G is a function of the temperature and CO₂ concentration, and therefore, indirectly, a function of the ventilation and enrichment rate. The optimal is found by computing the criterion over a fixed two-dimensional grid (ventilation rate and enrichment rate). The optimum control fluxes are then implemented in the ‘real’ greenhouse which is represented by the models described in Sections 4.1 and 4.2 (Figure 4). Due to the inaccuracy of the models used in determining the optimal control fluxes, the actual temperature and CO₂ concentration differ from the predicted ones. This, in turn, causes the actual criterion to be lower than the true optimum criterion which could be obtained if perfect greenhouse models were available.

4.1 Models of the ‘real’ greenhouse

The indoor air temperature is calculated according to

$$T_i(t + \Delta t) = (hT_s(t) + (\rho C Q(t) + U) T_o(t)) / (\rho C Q(t) + U + h). \quad (2)$$

In this equation, h is the air-soil heat transfer coefficient; T_s denotes the soil temperature; ρ and C are the air density and specific heat, respectively; Q is the ventilation rate; U is the cover heat transfer coefficient; T_o is the outdoor temperature; and Δt denotes the time-step (half an hour in the simulation). The first term on the right hand side denotes the heat exchange with the soil, in which heat storage takes place, and the second term denotes the sensible heat exchange

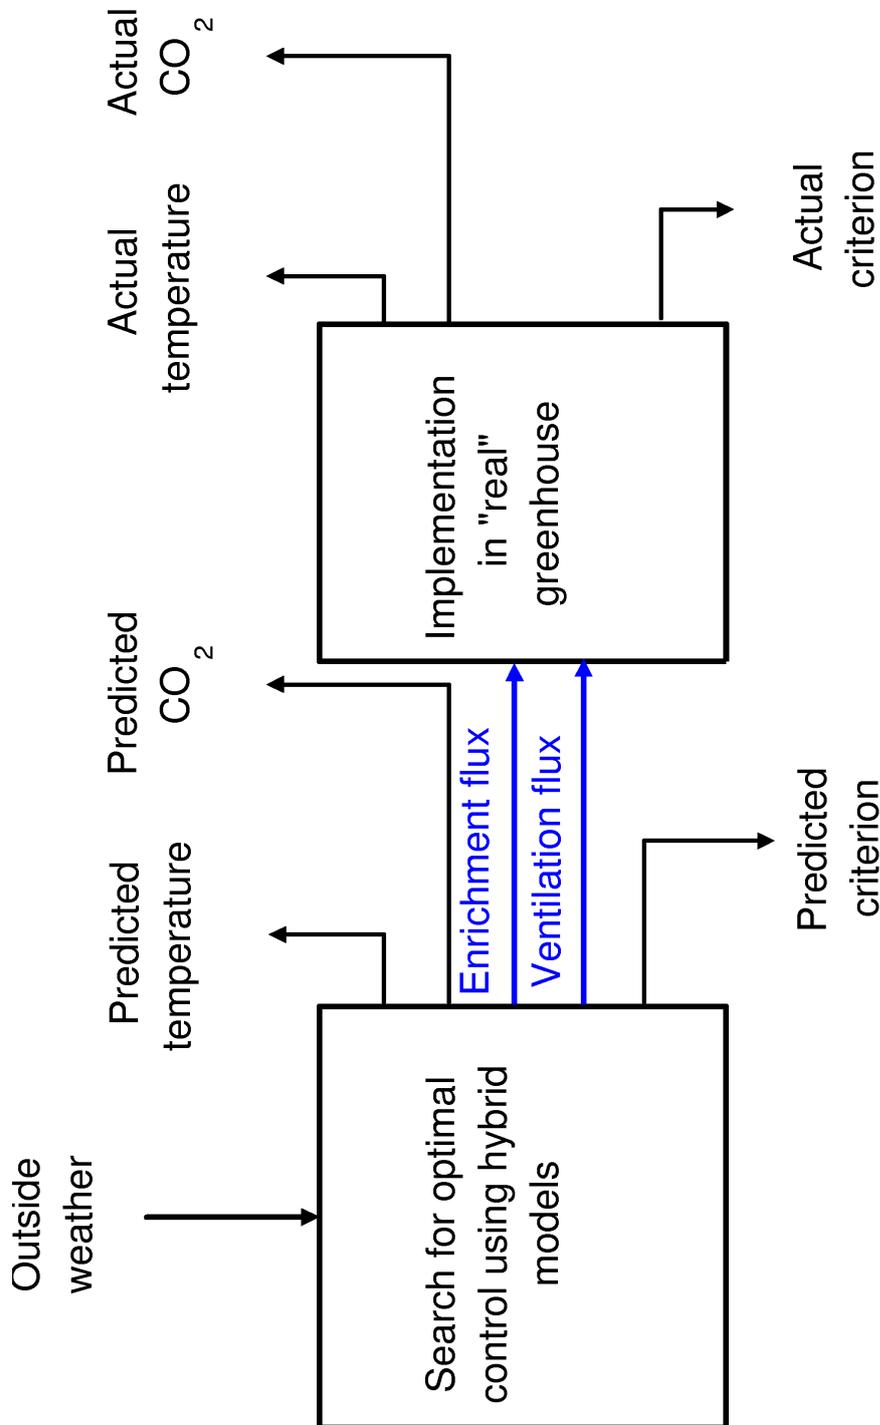


Figure 4: Schematic description of the simulation.

with the outdoor air through ventilation and conduction. Note that for simplicity, latent heat exchange is not modeled.

The soil temperature is described by:

$$T_s(t + \Delta t) = T_s(t) + (\lambda S_o(t) - h(T_s(t) - T_i(t)))/(Lc)\Delta t. \quad (3)$$

Here S_o is the solar radiation; λ is the solar radiation heating efficiency; L is the thickness of the soil layer in which the heat storage is assumed to take place; c is the soil heat capacity. This simplified model does not include the effect of long wave radiation, which is a dominant factor during nighttime. Therefore, cooling of the soil during nighttime is not simulated, but rather the initial morning temperature is reset daily to $T_o + 5$ [K], which is in good agreement with experimental observations.

The CO₂ concentration is described by:

$$X_i(t + \Delta t) = X_o + (R(t) - G(t))/(\rho Q(t)) \quad (4)$$

where X_i and X_o denote the indoor and outdoor CO₂ concentration, respectively; R denotes the enrichment rate, and G denotes the net photosynthetic flux, which is equal to the CO₂ uptake by the crop.

4.2 Crop model

The crop is described by

$$G = f_1 f_2 - f_3 \quad (5)$$

where

$$f_1 = \frac{\epsilon \zeta \tau S_o \gamma X_i}{(\epsilon \zeta \tau S_o + \gamma X_i)} \quad (6)$$

$$f_2 = \begin{cases} ((T_i - T_{min})^{0.5} (T_{max} - T_i)^{0.5}) / 11 & \text{if } T_i \leq T_{min} \text{ and } T_i \geq T_{max} \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

$$f_3 = w \phi e^{\nu(T_i - 25)} \quad (8)$$

f_1 denotes the Michaelis-Menten response to light CO₂ concentration, f_2 represent the effect of temperature on photosynthesis, and f_3 is the maintenance respiration rate. This model is similar to the one used by Linker *et al.* (1998b), except for the temperature function f_2 . This function f_2 has been adapted from Yin *et al.* (1995), and strongly penalizes temperatures close to T_{min} and T_{max} .

4.3 Greenhouse hybrid modeling

4.3.1 Physical models

The physical (manufacturer) model of the greenhouse air temperature is given by

$$T_i(t + \Delta t) = T_o(t) + 0.2S_o(t)/(\rho C Q(t) + 30) \quad (9)$$

This model includes only the effect of solar heating, and sensible heat exchange. In particular, heat storage in the greenhouse (in the soil) is not included in the model. This causes systematic

overestimation of the temperature during morning hours, and systematic underestimation during afternoon hours.

The physical model of the CO₂ concentration is

$$X_i(t + \Delta t) = X_o + R(t) / (\rho Q(t)). \quad (10)$$

This model is similar to the real greenhouse (Equation 4), except that it does not include the crop CO₂ uptake.

4.3.2 RBF models

The temperature RBF model has for inputs the ventilation rate, the outdoor temperature and solar radiation, and the soil temperature (all normalized between 0 and 1 using fixed values (Appendix C)). The CO₂ RBF model has for inputs the outdoor solar radiation, the ventilation rate, and the enrichment rate. These RBF models are adapted daily (at the end of the day), using databases formed with part of data collected up to that time (according to the method described in Section 3).

4.4 Results

Figure 5 presents average results for the whole simulation period (167 days), with different widths of the RBF activation functions σ , and different grid sizes (κ). The maximum number of datapoints associated with a center is five. The top frame of Figure 5 shows that the addition of the RBF models reduces the criterion loss from about 6 % (with the physical models alone) to about 1.5 % for κ between 0.2 and 0.5. The middle and bottom frames show that the temperature and CO₂ root mean square errors (RMSE) are also reduced, from about 2.5 [K] to 1 [K], and from 125 [ppm] to 35 [ppm], respectively. Very small values of κ lead to poorer results, because for such κ the activation function width σ is also very small, and the network has almost no extrapolation properties. Large values of κ also lead to poorer results, because the limited number of centers does not give enough flexibility to the network. All the results deteriorate when σ is reduced (relatively to κ), which corresponds to decreasing overlapping of adjacent activation functions.

The results obtained when the maximum number of datapoints per center is increased to ten or reduced to two (instead of five) are presented in Figures 6 and 7. Changing the maximum number of datapoints per center affects primarily the temperature RMSE. However, due to the weak influence of the accuracy of the temperature predictions on the performance criterion, the maximum number of datapoints per center has almost no effect on the performance criterion. As expected, the maximum number of datapoints per center influences strongly the size of the databases and, for the present application, a maximum of five datapoints per center is a good compromise between the size of the databases and the accuracy of the predictions.

Figure 8 shows the evolution of the number of centers for both models. At the beginning of the simulation, a large number of centers are added to correct the predictions of the physical models. For the CO₂ model, which has only three inputs (and has to model the CO₂ crop uptake), the number of centers stabilizes in about 10 days. For the temperature model, this process is longer, and stabilizes after about 25 days. However, as a result of changes in the outside weather, new centers are added between Days 160 and 180. Another change in the outside weather causes the addition of new centers during the last days of the simulation. Figure 9 presents the prediction errors (for the following day), in the particular case when $\kappa = \sigma = 0.3$. The error obtained if only the physical models are used is given for comparison.

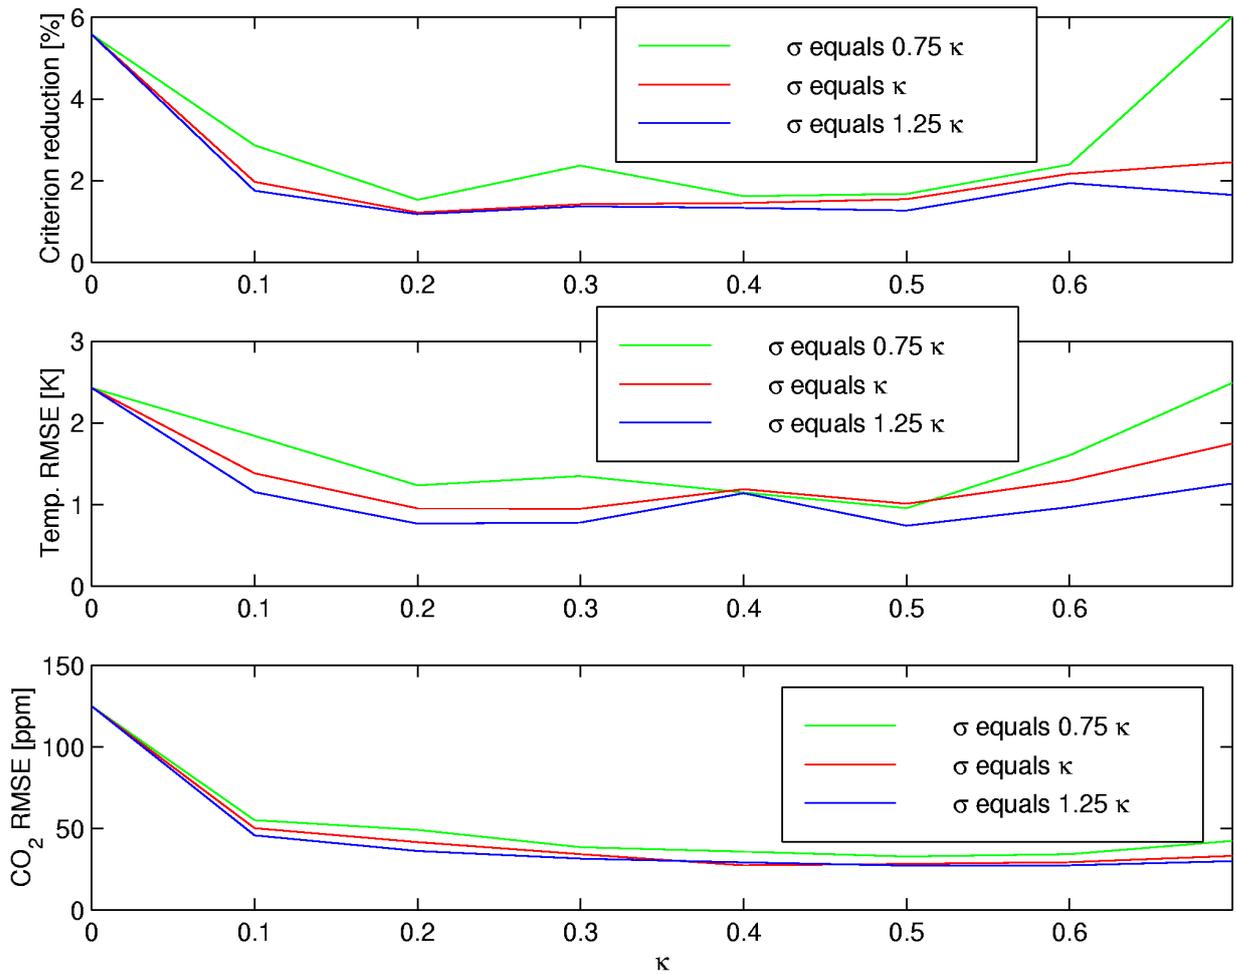


Figure 5: Criterion reduction (top frame), temperature RMSE (middle frame), and CO₂ RMSE (bottom frame), as function of the grid size κ . All the results are average results over the period considered. $\kappa = 0$ corresponds to the physical models alone. σ is the distance from a center at which the activation function equals 0.5.

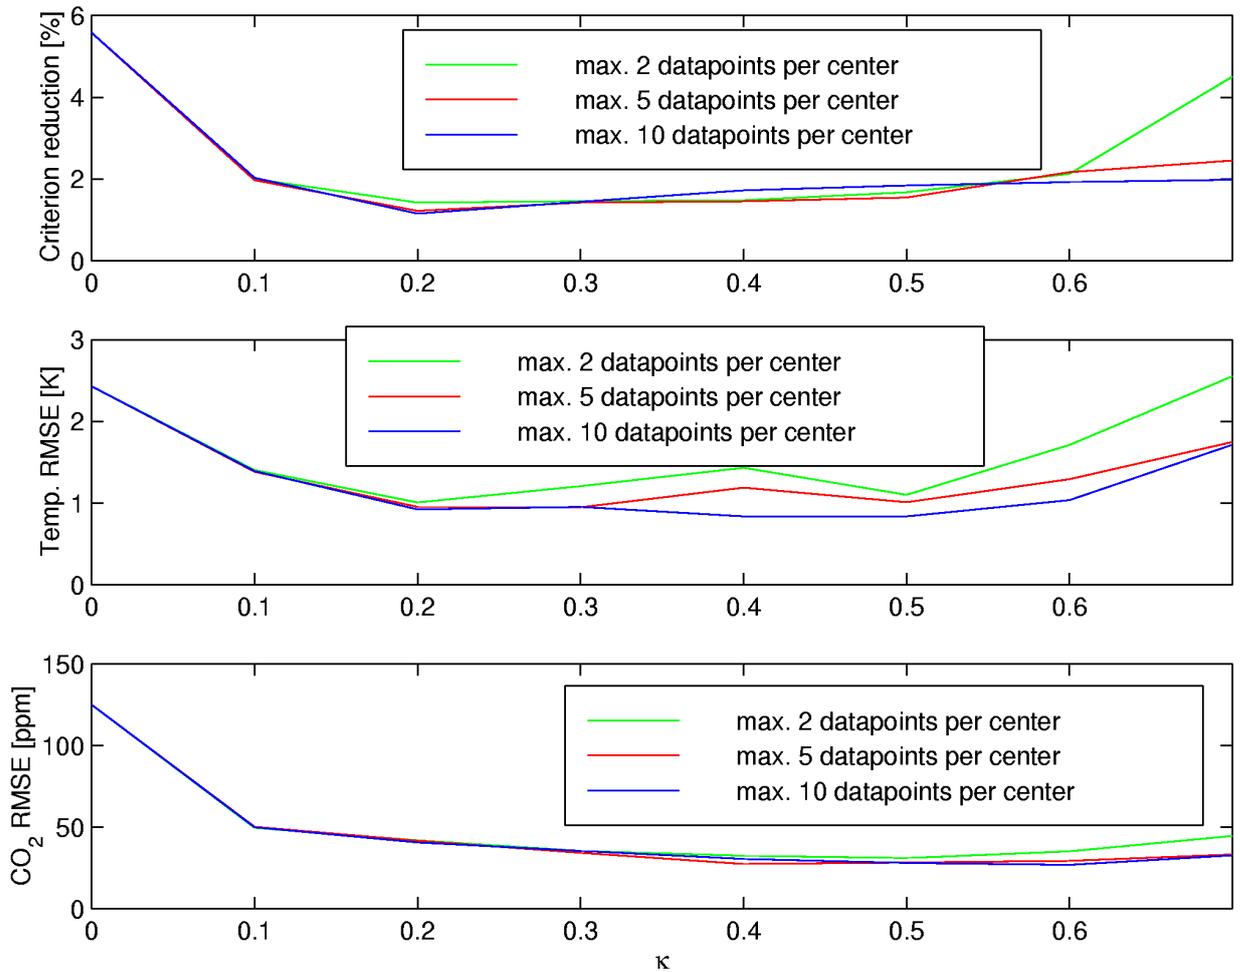


Figure 6: Criterion reduction (top frame), temperature RMSE (middle frame), and CO₂ RMSE (bottom frame), as function of the grid size κ . In all cases, $\sigma = \kappa$. $\kappa = 0$ corresponds to the physical models alone.

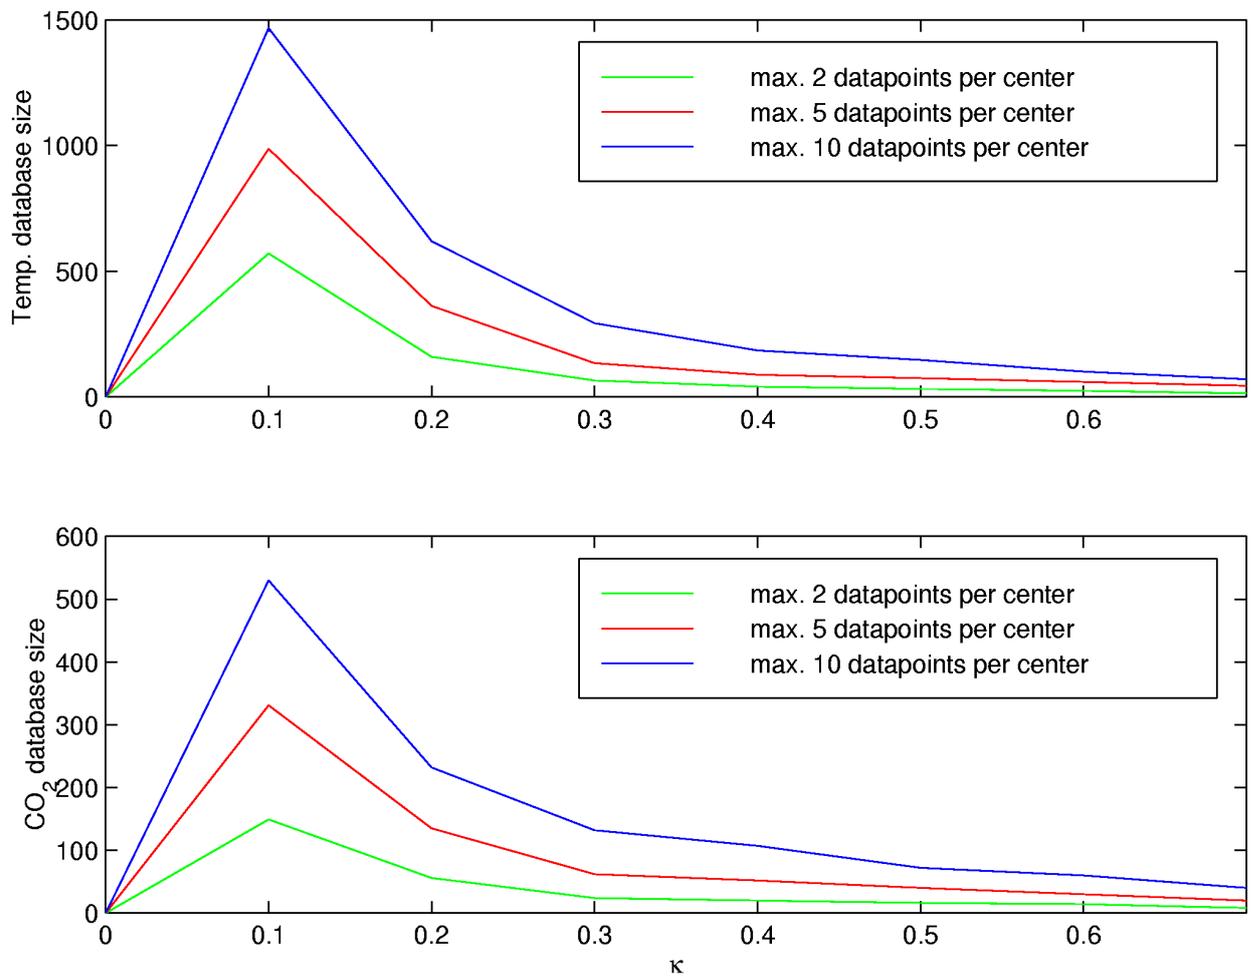


Figure 7: Size of the temperature (top frame) and CO₂ (bottom frame) database at the end of the simulation.

It can be seen that the prediction errors of the hybrid models quickly decrease, and comparison with Figure 8 shows that periods with relatively high RMSE correspond to periods during which new centers have to be added at the end of the day. The reduction of the daily criterion (relative to the true optimum criterion) is presented in Figure 10. After a short period of adaptation, the criterion calculated with the hybrid models becomes practically equal to the true optimum criterion which could be obtained if exact models of the system were known *a priori*.

5 Conclusions

Adaptive hybrid modeling offers the possibility to initiate a system's operation using an approximate model, and to use collected data to improve the model on-line. The adaptation is performed by enlarging the region over which a RBF network improves the predictions of a physical (or default) model which is approximately valid over the whole working domain. By comparison to standard model adaptation, in which the model is adapted only around the current working point, the training (or calibration) datapoints are selected so that they cover as large as possible a region of the input space.

Regarding the application to greenhouse climate modeling and optimization, the results show that, in the simplified situation considered (crop model known, no latent heat exchange, no heating required), the hybrid models adaptation takes place mostly during the first three weeks of operation. After this initial adaptation, new centers are added whenever the working point trajectory changes because of outside weather variations. By comparison with a situation in which fixed 'manufacturer' models would be used, the adaptive hybrid modeling improves the predictions' accuracy, and leads to a performance criterion very close to the true optimum criterion. This adaptive modeling approach has also been validated using experimental data (Linker *et al.*, 1998a; Linker, 1999), and more data are being collected to validate the method for a full growing season.

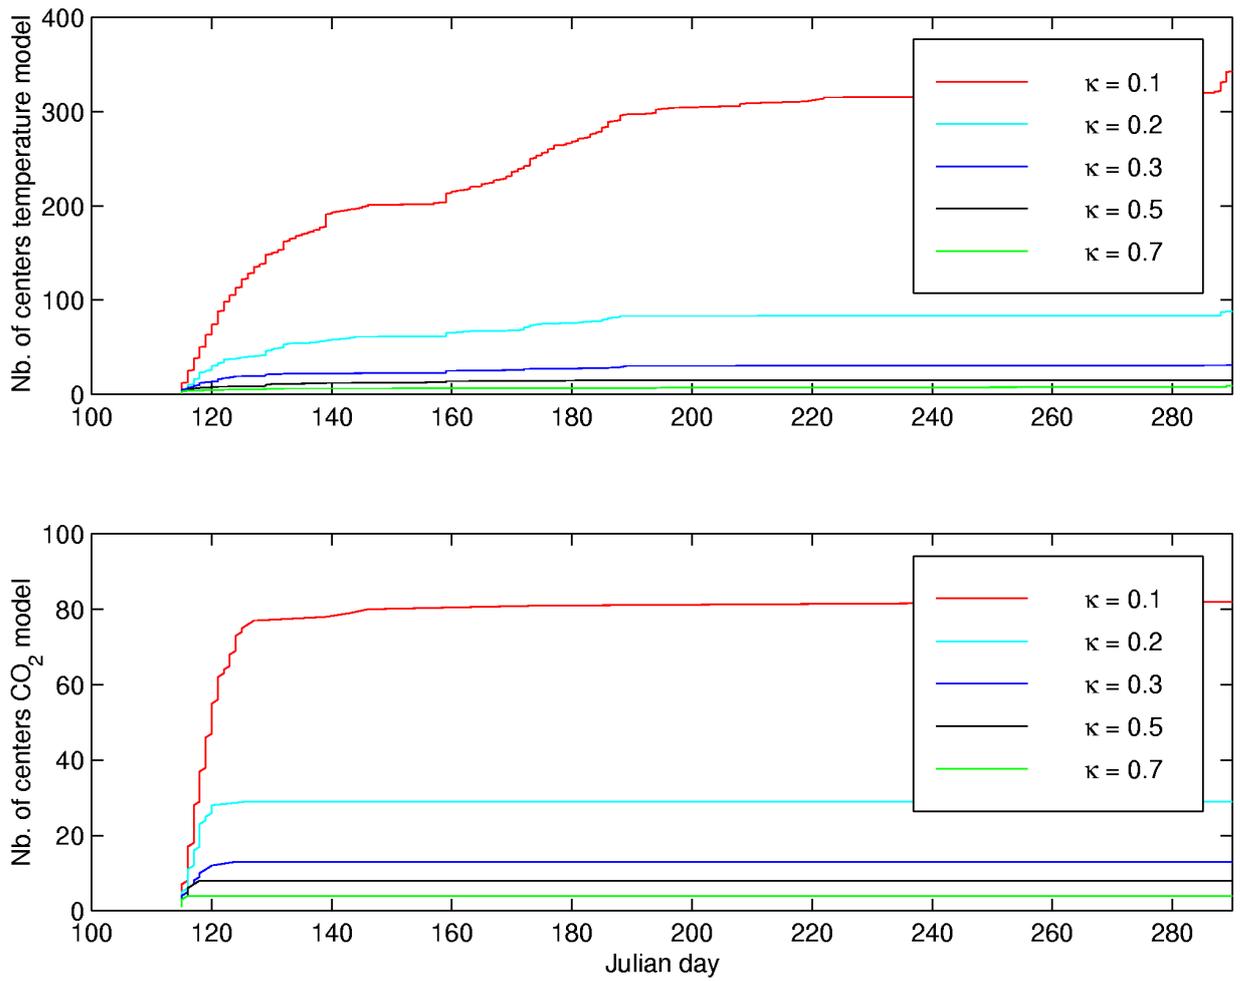


Figure 8: Number of centers of the temperature (top frame) and CO₂ (bottom frame) RBF models, as a function of the julian day. Day 115 corresponds to April 25.

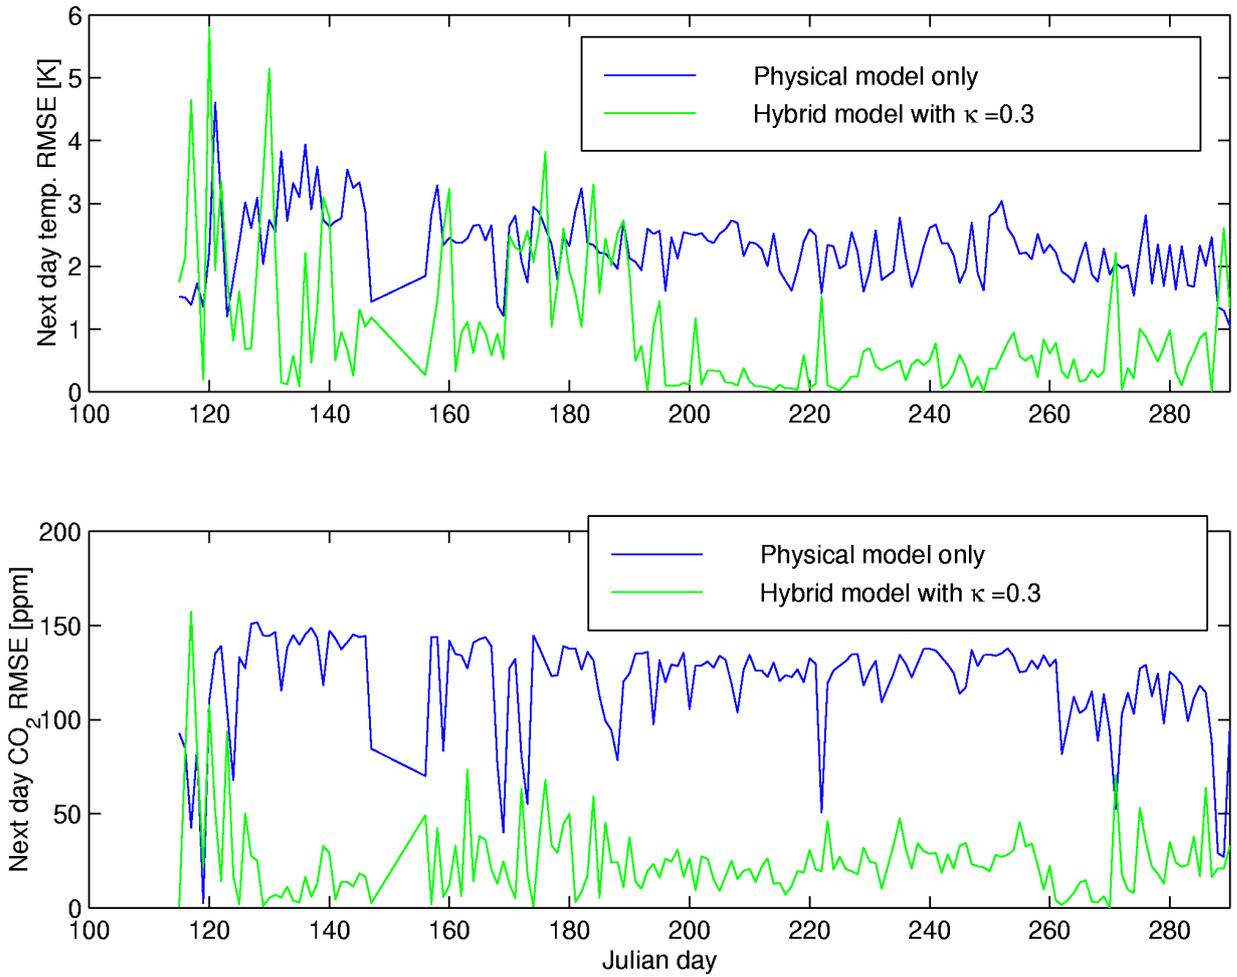


Figure 9: RMSE of the temperature (top frame) and CO₂ (bottom frame) models. The predictions for Day $d + 1$ are calculated with the hybrid models trained with part of the data collected up to Day d . RMSE of the physical (fixed) models are given for reference.

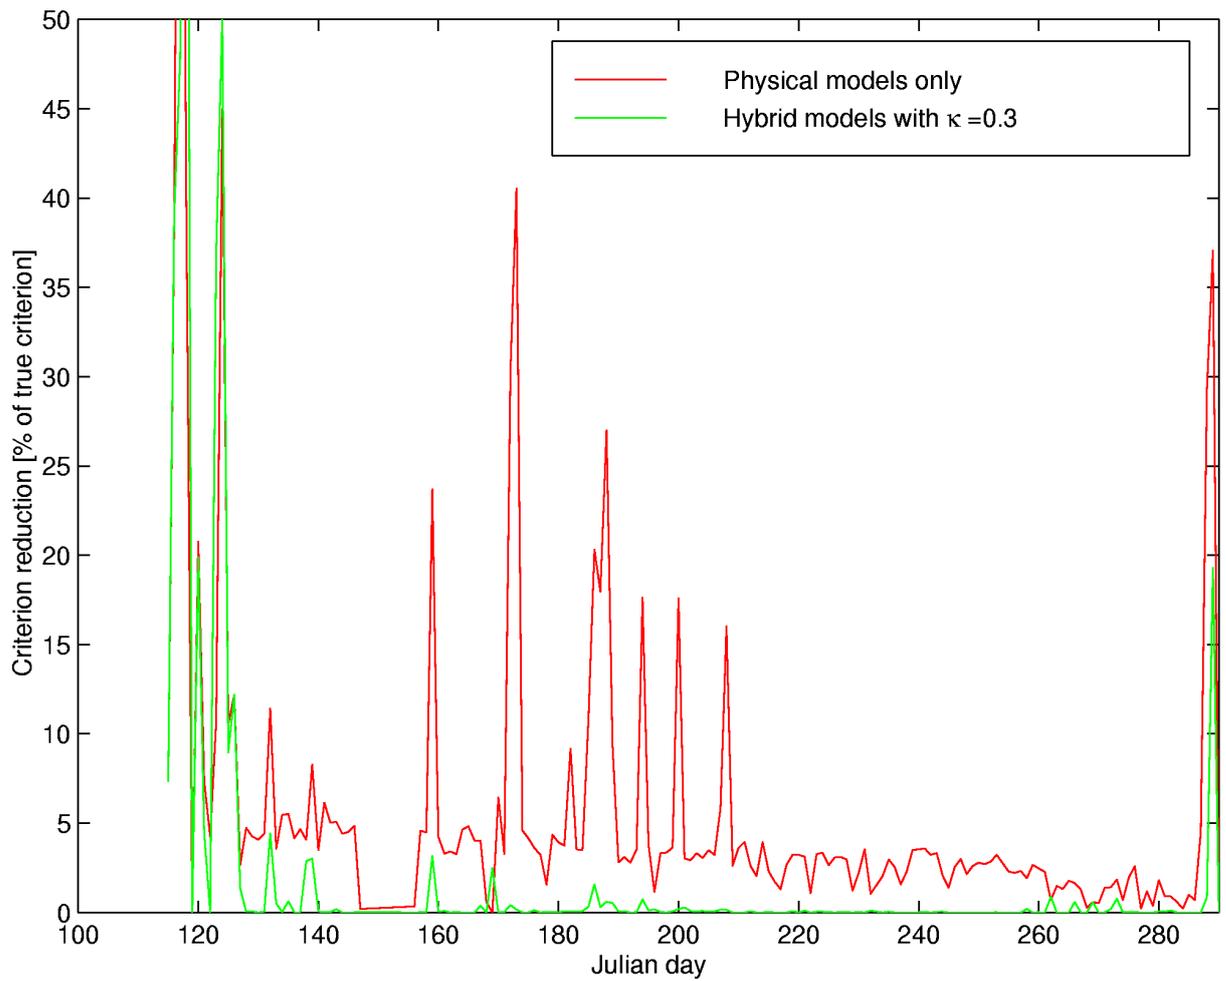


Figure 10: Daily criterion reduction, expressed in percent of the true optimum criterion.

A Main symbols

C	Specific heat of air at constant pressure	$J/(kg(air)K)$
c	Soil heat capacity	$J/(m^3(ground)K)$
G	Net photosynthesis flux	$kg(CO_2)/(m^2(ground)s)$
h	Air-soil heat transfer coefficient	$W/(m^2(ground)K)$
j	Instantaneous performance criterion	$\$/m^2(ground)s$
K_c	Unit price of supplied CO_2	$\$/kg(CO_2)$
K_g	Unit price of CO_2 in biomass	$\$/kg(crop CO_2)$
L	Soil layer thickness	$m(ground)$
Q	Ventilation rate	$m^3(air)/(m^2(ground)s)$
R	Enrichment rate	$kg(CO_2)/(m^2(ground)s)$
S	Solar radiation	$W/m^2(ground)$
T	Temperature	K
t	Time	s
U	Heat transfer coefficient of greenhouse cover	$W/(m^2(ground)K)$
w	CO_2 content of crop	$kg(crop CO_2)/(m^2(ground))$
X	CO_2 concentration of the air	$kg(CO_2)/kg(air)$
γ	Leaf conductance to CO_2	m/s
ϵ	Photosynthetic efficiency	$kg(CO_2)/J(PAR)$
ζ	Ratio of PAR to solar radiation	-
κ	Radial basis function grid size	-
λ	Solar radiation heating efficiency	-
ν	Respiration exponent	K^{-1}
ρ	Air density	$kg(air)/m^3(air)$
σ	Distance from RBF center at which the activation function equals 0.5	-
τ	Cover transmissivity	-
ϕ	Respiration rate per unit crop mass	$kg(CO_2)/(kg(crop CO_2)s)$

Subscripts

i	indoor
o	outdoor
s	soil

Acronyms

NN	Neural Network
PAR	Photosynthetic Active Radiation
RBF	Radial Basis Function
RMSE	Root Mean Square Error

B Parameter values

B.1 Greenhouse models

$$\begin{aligned}
 U &= 10 \text{ [W/m}^2\text{]} \\
 c &= 2912 \times 10^3 \text{ [J/(m}^3\text{K)}] \\
 h &= 20 \text{ [W/(m}^2\text{K)}] \\
 L &= 0.2 \text{ [m]} \\
 \lambda &= 0.5 \text{ [-]}
 \end{aligned}$$

B.2 Crop model

$$\begin{aligned}
 T_{min} &= 15 \text{ [}^\circ\text{C]} \\
 T_{max} &= 37 \text{ [}^\circ\text{C]} \\
 w &= 0.1 \text{ [kg(crop CO}_2\text{)/m}^2\text{(ground)}] \\
 \epsilon &= 10^{-8} \text{ [kg(CO}_2\text{)/J(PAR)}] \\
 \gamma &= 2 \times 10^{-3} \text{ [m/s]} \\
 \nu &= 0.0693 \text{ [K}^{-1}\text{]} \\
 \phi &= 0.4 \times 10^{-6} \text{ [kg(CO}_2\text{)/(kg(crop CO}_2\text{)s)}] \\
 \tau &= 0.5 \text{ [-]} \\
 \zeta &= 0.7 \text{ [-]}
 \end{aligned}$$

B.3 Prices

$$\begin{aligned}
 K_c &= 0.4 \text{ [\$ / (kg(CO}_2\text{))]} \\
 K_g &= 15 \text{ [\$ / (kg(CO}_2\text{))]}
 \end{aligned}$$

C Normalization values

Each variable is normalized according to:

normalized value = (value - minimum value) / (maximum value - minimum value).

Variable	Minimum	Maximum	Units
CO ₂ RBF output	-300×10^{-6}	300×10^{-6}	kg(CO ₂)/kg(air)
Enrichment rate	0	40×10^{-6}	kg(CO ₂)/(m ² s)
Soil temperature	10	45	°C
Indoor temperature	10	45	°C
Outdoor solar radiation	0	1000	W/m ²
Outdoor temperature	10	35	°C
Temperature RBF output	-8	8	°C
Ventilation rate	0	0.05	m ³ /(m ² s)

References

- Avissar, R. and Y. Mahrer (1982). "Verification study of a numerical greenhouse microclimate model," *Transactions of the ASAE*, **25**, pp. 1711-1720.
- Bishop, C. M. (1993). "Curvature-driven smoothing: A learning algorithm for feedforward networks," *IEEE Transactions on Neural Networks*, **4**, no. 5, pp. 882-884.

- Boulard, T. and A. Baille (1993). "A simple greenhouse climate control model incorporating effects of ventilation and evaporative cooling," *Agricultural and Forest Meteorology*, **65**, pp. 145–157.
- Chalabi, Z. S. and W. Zhou (1997). "Optimization of CO₂ enrichment in greenhouses constrained by limited CO₂ supply and CO₂ dynamics," in *Preprints of the 3rd IFAC/ISHS Workshop on Mathematical and Control Applications in Agriculture* (A. Munack and H. J. Tantau, eds.), pp. 109–112.
- Chen, S., S. A. Billings, and P. M. Grant (1992). "Recursive hybrid algorithm for non-linear system identification using radial basis function networks," *International Journal of Control*, **55**, no. 5, pp. 1051–1070.
- Feyo de Azevedo, S., B. Dahm, and F. R. Oliveira (1997). "Hybrid modelling of biochemical processes: A comparison with the conventional approach," *Computers and Chemical Engineering*, **21**, no. Suppl., pp. s751–s756.
- Gutman, P. O., P. O. Lindberg, I. Ioslovich, and I. Seginer (1993). "A non-linear optimal greenhouse control problem solved by linear programming," *Journal of Agriculture Engineering Research*, **55**, pp. 335–351.
- Ioslovich, I., I. Seginer, P. O. Gutman, and M. Borshchevsky (1995). "Sub-optimal CO₂ enrichment of greenhouses," *Journal of Agriculture Engineering Research*, **60**, pp. 117–136.
- Joerding, W. H. and J. L. Meador (1991). "Encoding a priori information in feedforward neural networks," *Neural Networks*, **4**, pp. 847–856.
- Jordan, M. I. and D. E. Rumelhart (1992). "Forward models: supervised learning with a distal teacher," *Cognitive Science*, **16**, pp. 307–354.
- Kindelan, M. (1980). "Dynamic modelling of greenhouse environment," *Transactions of the ASAE*, **23(5)**, pp. 1232–1239.
- Kok, R., R. Lacroix, G. Clark, and E. Taillefer (1994). "Imitation of a procedural greenhouse model with an artificial neural network," *Canadian Agricultural Engineering*, **36(2)**, pp. 117–126.
- Linker, R. (1999). *Failure detection and isolation in greenhouses*, Ph.D. thesis, Technion, Israel Institute of Technology. Haifa, Israel. PostScript file available upon request at: linkerr@tx.technion.ac.il.
- Linker, R., I. Seginer, and P. O. Gutman (1998a). "Neural network and hybrid adaptive modeling of greenhouse air temperature," in *Proceedings of the AgEng'98 conference (in press)*.
- Linker, R., I. Seginer, and P. O. Gutman (1998b). "Optimal control of CO₂ in a greenhouse modelled with neural networks," *Computers and Electronics in Agriculture*, **19**, pp. 289–310.
- Milanic, S., S. Sel, N. Hvala, S. Strmcnik, and R. Karba (1997). "Applying artificial neural network models to control a time variant chemical plant," *Computers and Chemical Engineering*, **21**, no. Suppl., pp. s637–s642.
- Moody, J. and C. J. Darken (1989). "Fast learning in networks of locally-tuned processing units," *Neural Computation*, **1**, pp. 281–294.
- Piron, E., E. Latrille, and F. Rene (1997). "Application of artificial neural networks for crossflow micro-filtration modeling: "blac-box" and semi-physical approaches," *Computers and Chemical Engineering*, **21**, no. 9, pp. 1021–1030.
- Platt, J. (1991). "A resource-allocating network for function interpolation," *Neural Computation*, **3**, pp. 213–225.
- Psichogios, D. C. and L. H. Ungar (1992). "A hybrid neural network-first principles approach to process modeling," *American Institute of Chemical Engineers Journal*, **38**, no. 10, pp. 1499–1511.

- Rudolph, S. (1997). "On topology, size and generalization of non-linear feed-forward neural networks," *Neurocomputing*, **16**, pp. 1–22.
- Seginer, I., A. Angel, S. Gal, and D. Kantz (1986). "Optimal CO₂ enrichment strategy for greenhouses: a simulation study," *Journal of Agriculture Engineering Research*, **34**, pp. 285–304.
- Seginer, I., T. Boulard, and B. J. Bailey (1994). "Neural network models of the greenhouse climate," *Journal of Agriculture Engineering Research*, **59**, pp. 203–216.
- Takakura, T., K. A. Jordan, and L. L. Boyd (1971). "Dynamic simulation of plant growth and environment," *Transactions of the ASAE*, **14**, pp. 964–971.
- Tanomaru, J. and S. Omatu (1992). "Process control by on-line trained neural controllers," *IEEE Transactions on Industrial Electronics*, **39**, no. 6, pp. 511–521.
- Thompson, M. L. and M. A. Kramer (1994). "Modelling chemical processes using prior knowledge and neural networks," *American Institute of Chemical Engineers Journal*, **40**, no. 8, pp. 1328–1340.
- Tsen, A. Y., S. S. Jang, and D. S. H. Wong (1996). "Predictive control of quality in batch polymerization using hybrid ANN models," *American Institute of Chemical Engineers Journal*, **42**, no. 2, pp. 455–465.
- Wilson, J. A. and L. F. M. Zorretto (1997). "A generalized approach to process state estimation using hybrid artificial neural network/mechanistic models," *Computers and Chemical Engineering*, **21**, no. 9, pp. 951–963.
- Yin, X., M. J. Kropff, G. McLaren, and R. M. Visperas (1995). "A nonlinear model for crop development as a function of temperature," *Agricultural and forest meteorology*, **77**, pp. 1–16.
- Yingwei, L., N. Sundararajan, and P. Saratchandran (1997). "Identification of time-varying nonlinear systems using minimal radial basis function neural networks," *IEE Proceedings D-Control Theory and Applications*, **144**, no. 2, pp. 202–208.