Computers and Chemical Engineering 49 (2013) 37–49

Contents lists available at SciVerse ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Dynamic optimization of bioreactors using probabilistic tendency models and Bayesian active learning

Ernesto C. Martínez[b,*], Mariano D. Cristaldi[b], Ricardo J. Grau[a]

[a] INTEC (CONICET-UNL), Güemes 3450, Santa Fe 3000, Argentina
[b] INGAR (CONICET-UTN), Avellaneda 3657, Santa Fe S3002 GJC, Argentina

ARTICLE INFO

Article history:
Received 24 August 2011
Received in revised form 17 August 2012
Accepted 20 September 2012
Available online 16 October 2012

Keywords:
Bayesian inference
Experimental design
Fed-batch fermentation
Modeling for optimization
Run-to-run optimization
Sensitivity analysis

ABSTRACT

Due to the complexity of metabolic regulation, first-principles models of bioreactor dynamics typically have built-in errors (structural and parametric uncertainty) which give rise to the need for obtaining relevant data through experimental design in *modeling for optimization*. A run-to-run optimization strategy which integrates imperfect models with Bayesian active learning is proposed. Parameter distributions in a probabilistic model of bioreactor performance are re-estimated using data from experiments designed for maximizing information and performance. The proposed Bayesian decision-theoretic approach resorts to probabilistic tendency models that explicitly characterize their levels of confidence. Bootstrapping of parameter distributions is used to represent parametric uncertainty as histograms. The Bajpai & Reuss bioreactor model for penicillin production validated with industrial data is used as a representative case study. Run-to-run convergence to an improved policy is fast despite significant modeling errors as long as data are used to revise iteratively posterior distributions of the most influencing model parameters.

1. Introduction

Most bioprocess optimization techniques are model-based (De Tremblay, Perrier, Chavarie, & Archambault, 1993; Frahm, Lane, Märk, & Pörtner, 2003; Guthke & Knorre, 1981; Lim, Tayeb, Modak, & Bonte, 1986; Riascos & Pinto, 2004), and since accurate models are rarely available, experimental optimization of the operating policy is a difficult problem to be addressed for a successful scale-up. The best use of an imperfect first-principles model through proper handling of its inherent uncertainty is a challenging issue for fast productivity improvement of innovative fed-batch fermentations using data sampled from a small number of production runs. Bioreactors are engineered systems in which the activity of living cells is harnessed to produce an antibiotic, antibody, protein, a tissue or a host of other products of considerable impact on human life (Anesiadis, Cluett, & Mahadevan, 2008; Jain & Kumar, 2008; Ramkrishna, 2003). For maximum productivity, cells in a bioreactor must be maintained in an appropriate state of metabolic activity by tightly controlling conditions in the abiotic phase. The main problem in bioreactor modeling for optimization is that biological activity occurs in metabolic pathways which are controlled by switches through built-in regulatory networks (Geng & Yuan, 2010). Due to the complexity of metabolic regulation and limited measurements, first-principles models of bioreactor dynamics can only capture the qualitative tendency of sampled state variables such as biomass, substrate and product concentrations (Martínez, Cristaldi, & Grau, 2009; Tsobanakis, 1994). Hence, without biasing data gathering by increasingly improving the operating policy, bioreactor performance predictions are too uncertain and unreliable in quantitative terms to be useful for productivity optimization (Bonvin, 1998; Martínez & Wilson, 2003; Schenker & Agarwal, 1995). As a result, migration from laboratory conditions to production runs is often made with high levels of uncertainty about the degree of optimality of an operating policy (Terwiesch, 1995; Terwiesch & Agarwal, 1995). Consequently, a very conservative and sub-optimal operating policy is repeatedly applied to industrial bioreactors seeking reproducibility rather than improvement (Martínez & Wilson, 2003).

Run-to-run optimization of the operating policy for a fed-batch bioreactor using data gathered in production runs can be approached using two alternatives: (i) a systematic model-based iteration strategy, or (ii) a heuristic procedure using somehow past operating experience for modifying the policy directly. The heuristic optimization approach based on intuitively tweaking input profiles is very inefficient, often leads to sub-optimal solutions, and it cannot guarantee neither systematic performance improvement nor convergence to a near optimal policy. An interesting step in this direction has been proposed in Smets, Claes, November, Bastin, and Van Impe (2004) by starting from a model-derived operating policy and optimal profiles of the key state variables. Then, the optimal

* Corresponding author. Tel.: +54 342 4534451; fax: +54 342 4553439.
E-mail address: ecmarti@santafe-conicet.gov.ar (E.C. Martínez).

0098-1354/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compchemeng.2012.09.010

Nomenclature

F_{in}	Inlet flow rate ($L\,h^{-1}$)
F_{evap}	outlet flow rate due evaporation ($L\,h^{-1}$)
J	performance index
$\boldsymbol{m}(t)$	time-dependent control variables.
$p(\theta_i)$	*prior* distribution of the *i*th model parameter.
$p(\boldsymbol{\theta}\|\boldsymbol{x})$	*posterior* probability distribution for parameters.
P	penicillin concentration (as potassium salt) (g PenGK L^{-1})
Q	global sensitivity matrix
S	substrate concentration ($g\,L^{-1}$)
t	time (h)
t_f	final time of an experimental run (h)
$\mathbf{t}^{sp}$	vector of sampling times in an evaluation experiment
V	culture broth volume (L)
$\boldsymbol{x}(t)$	vector of state variables
$\boldsymbol{x}(t_i)$	bioreactor sampled state vector at a given time.
X	biomass concentration (g-DW L^{-1})
X_d	death biomass concentration (g-DW L^{-1})
X_v	viable biomass concentration (g-DW L^{-1})
u	utility function
$\boldsymbol{w}$	time-invariant control variables

Parameters

A	feeding profile parameter ($L\,h^{-2}$)
B	feeding profile parameter (h^{-1})
C	feeding profile parameter (h^{-2})
T_{feed}	initial time for fed-batch operation (h)

Greek symbols

Θ	feasible set of model parameters
$\boldsymbol{\beta}$	set of parameters describing time-varying inputs
$\boldsymbol{\varphi}$	vector of operating policy parameters
μ	specific biomass growth rate (h^{-1})

solution is implemented in the form of a model-independent sub-optimal strategy by using a modified (semi-empirical) control function, which includes reduced terms based on heuristic observations. More effective, though, is designing dynamic experiments to extract useful information from policy evaluation runs. In this way, the operating policy is improved by introducing relevant data for optimization in an imperfect model. This approach does not rely on expert knowledge, but requires to model available data carefully. For model-based policy optimization to be successful it is mandatory to re-estimate selectively the more sensitive model parameters using optimal experimental design techniques in data gathering (Martínez et al., 2009).

An approach for model-based heuristic optimization of operating policies has been proposed in Maria (2004, 2007) and successfully applied to D-glucose oxidation. This author argues that, by using reduced order (low complexity) bioreactor models and through semi-empirical optimal control functions, it is possible to lower computational costs and experimental efforts necessary to identify and verify all model parameters and reaction steps under a wide range of operating conditions and at different time scales. The reduced order model is based on a simplified enzymatic kinetics, requires a small number of on-line measurements for model update and a few parameters are used to adjust the control function. The solution found is implemented in the form of a model-independent sub-optimal strategy based on a control function selected from a library. However, the heuristic optimization approach is highly problem-dependent (*e.g.*, enzyme oxidation) since it mostly relies on an intricate understanding of the characteristics of the bioprocess behavior and human judgment for defining an improved policy while addressing the dilemma of knowledge exploitation *versus* exploring untried operating conditions. This dilemma is at the very heart of modeling for optimization with imperfect models. When a reduced order model is used for policy improvement you cannot improve its parametric precision comprehensively. Thus, the model is only a means to find better policies at the cost of biasing data gathering in the most profitable region of operating conditions. Lacking a conceptual framework for policy optimization, generalization and incorporation of uncertainty into the decision-making process, the heuristic optimization approach is costly in terms of both time and money. Expert knowledge can be difficult to obtain, expensive, or is simply not available. Moreover, no systematic reduction of model uncertainty is made as more experimental data is available which prevents guaranteeing steady policy improvement and convergence toward a near-optimal solution.

In the attempt to compensate for a significant process-model mismatch, optimal operation under uncertainty requires using measurements from carefully designed experiments to improve on a run-to-run basis from a cautious (sub-optimal) policy. This model-based policy optimization approach consists of iteratively using new measurements to increasingly reduce parametric uncertainty in a tendency (imperfect) model and later resorting to the updated model for policy improvement (Martínez et al., 2009). A "tendency model" is a low order, nonlinear, dynamic model that approximates the stoichiometry and kinetic relationships of a bioprocess using the available plant data along with fundamental knowledge of the process characteristics (Bonvin & Rippin, 1990; Filippi, Bordet, Villermaux, Marchal-Brassey, & Georgakis, 1989; Fotopoulos, Georgakis, & Stenger, 1998; Georgakis, 1995; Uhlemann, Cabassud, LeLann, Borredon, & Cassamatta, 1994). Operating policies based on over-confident first-principles models often fail to yield productivity improvement due to a lack of parametric precision and structural errors.

For Bayesian optimization with tendency models, not only a bioreactor model for policy improvement is required, but it is also important that the model faithfully describes its own accuracy to treat uncertainties in a principled way. Humans do something similar: as it is argued in (Körding & Wolpert, 2004, 2006), whenever humans have only little experience, they employ an internal forward model for predictions and average over the uncertainty when extrapolating and making decisions. The essential characteristic of Bayesian methods is their explicit use of probability theory for quantifying uncertainty in inferences based on statistical data analysis. Without any notion of uncertainty, the model-optimized policy would be too confident and claims exact knowledge, which it actually does not have. Representation and incorporation of model uncertainty in run-to-run optimization is particularly important in the early stages of bioprocess scale-up when the available data set is very sparse and has been obtained for a wide range of operating conditions. For Bayesian optimization of bioreactors, the novel concept of a *probabilistic tendency model* that integrates first-principles and constitutive laws with probability distributions for describing parametric uncertainty is proposed.

In this work, a general and fully Bayesian decision-theoretic framework for policy optimization in innovative bioprocesses is presented. In the case of only few production runs with a full-scale bioreactor, the problem of dealing with fairly limited experience to improve the policy is successfully addressed using Bayesian active learning. In Bayesian inference, scarce experimental data are used to learn a probabilistic model of a bioreactor dynamics by updating parameter distributions. Probabilistic tendency models are able to represent and to quantify their own uncertainty for safe generalization of available experience to untried operating conditions. Thus, uncertainty is explicitly accounted for in run-to-run