

Optimized Treatment of Fibromyalgia Using System Identification and Hybrid Model Predictive Control

Sunil Deshpande^a, Naresh N. Nandola^{a,1}, Daniel E. Rivera^{a,*}, Jarred W. Younger^b

^a*Control Systems Engineering Laboratory, School for Engineering of Matter, Transport, and Energy, Arizona State University, Tempe, AZ 85287 USA.*

^b*Department of Psychology, Departments of Anesthesiology and Rheumatology, University of Alabama at Birmingham, Birmingham, AL 35233 USA.*

Abstract

The term *adaptive intervention* is used in behavioral health to describe individually-tailored strategies for preventing and treating chronic, relapsing disorders. This paper describes a system identification approach for developing dynamical models from clinical data, and subsequently, a hybrid model predictive control scheme for assigning dosages of naltrexone as treatment for fibromyalgia, a chronic pain condition. A simulation study that includes conditions of significant plant-model mismatch demonstrates the benefits of hybrid predictive control as a decision framework for optimized adaptive interventions. This work provides insights on the design of novel personalized interventions for chronic pain and related conditions in behavioral health.

Keywords: optimized adaptive behavioral interventions, fibromyalgia, system identification, hybrid model predictive control, biomedical applications

1. Introduction

With rising health care costs, there is an increasing interest in the medical community towards developing improved strategies for treating chronic diseases (Wellstead et al., 2008; Collins, 2010). Among these lie adaptive interventions, which consider adjusting treatment dosages over time based on participant response. Control engineering offers a broad-based solution framework for optimizing the effectiveness of such interventions and has been proposed as an enabler for more efficacious treatments that minimize waste, increase compliance, and enhance the intervention potency (Rivera et al., 2007; Zafra-Cabeza et al., 2011; Riley et al., 2011; Deshpande et al., 2014).

Traditional medical practice is based on treatment protocols designed for a standard response that do not necessarily incorporate individual characteristics or optimization procedures. Many of these dosage strategies are aimed at acute disorders and in spite of effective drugs, are not

*Corresponding author. Tel.: +1-(480)965-9476. Fax: +1-(480)727-9321

Email address: daniel.rivera@asu.edu (Daniel E. Rivera)

¹Currently at ABB Corporate Research Center, Bangalore, India.

necessarily efficient for relapsing, chronic disorders. The use of adaptive approaches, in which dosages are adjusted based on participant response over time, is the key motivation for use of control systems engineering principles. This paper demonstrates how control engineering principles can impact the treatment of chronic, relapsing disorders by examining a pain condition known as fibromyalgia (FM) (Boissevain & McCain, 1991a,b; Younger & Mackey, 2009; Deshpande et al., 2011). The examination is based on secondary analysis of information collected from a previously conducted clinical trial using naltrexone for the treatment of FM. This problem is approached from a systems and controls point-of-view: first, system identification techniques are applied to develop dynamical models from daily diary reports completed by intervention participants. These diary reports include self-assessments of outcomes of interest (e.g., general pain symptoms, sleep quality) and additional external variables that affect these outcomes (e.g., stress, anxiety, and mood). These dynamical system models serve as the basis for applying model predictive control as a decision algorithm for dosage selection of naltrexone. The categorical/discrete-event nature of the dosage assignment process calls for hybrid model predictive control (HMPC) schemes. Instead of relying on conventional tuning of HMPC using weight matrices, a multiple degree-of-freedom formulation is evaluated in this paper that enables the user to adjust the speed of setpoint tracking, measured disturbance rejection and unmeasured disturbance rejection independently in the closed loop system. Simulation results depicting realistic conditions are presented to illustrate the benefits of the proposed control scheme in addressing hybrid dynamics, clinical constraints and plant-model mismatch typically present in such applications.

The paper is organized according to the following sections: Section 2 briefly describes the intervention and nature of the associated clinical data. Section 3 discusses the procedure for building parsimonious models using system identification. The HMPC formulation used for dosage assignment is presented in Section 4, with Section 5 demonstrating the application of HMPC for delivering adaptive interventions under conditions of disturbances and model uncertainty. The paper ends with a summary and conclusions in Section 6.

2. Naltrexone intervention for fibromyalgia

Fibromyalgia (FM) is a disorder characterized primarily by chronic widespread pain. The characteristic symptoms of FM are diffuse musculoskeletal pain and sensitivity to mechanical stim-

ulation at soft tissue tender points (Wolfe et al., 1990, 2010). Other important symptoms of FM include fatigue, sleep irregularities, bowel abnormalities, anxiety, and mood dysfunction. While no specific laboratory test can confirm FM, most patients present with a history of widespread pain and fatigue conditions. Another important issue with FM is that its etiology is largely unknown and without any scientific consensus (Perrot, 2008), although the condition is suspected to involve central sensitization of pain processing (Lee et al., 2011). As the causes of FM are unknown, it has been difficult to single out a specific type of treatment for this chronic disease. Depending on different approaches for the mechanisms of FM, there have been experiments with various drugs. There is a good evidence to suggest that naltrexone, an opioid antagonist, has a neuroprotective role and may be a potentially effective treatment for diseases like FM (Younger & Mackey, 2009; Mattioli et al., 2010). The data for this paper has been taken from clinical trials of a low dose naltrexone (LDN) intervention conducted by Dr. Jarred Younger and colleagues at the Stanford Systems Neuroscience and Pain Lab (SNAPL), Stanford University School of Medicine (Younger & Mackey, 2009; Younger et al., 2013).

2.1. The data

The study was conducted in two phases: a single blind pilot study on 10 participants and a double blind full study on 30 participants; the latter also featured a protocol of longer duration. A crossover design was employed where participants received both drug and placebo, and hence act as their own control. A fixed naltrexone dose of concentration 4.5 mg was administered. In the pilot study, the participants received placebo followed by drug (P-D protocol) whereas the full study participants were randomized to receive either drug first (D-P protocol) or placebo first (P-D protocol). The time line is split into baseline (during this phase participants do not receive any kind of medication), placebo/drug and finally washout phase (all kinds of medications are stopped). The number of data points ranged from 98 to 154 at daily sampling ($T = 1$). Participants entered their responses in a handheld computer to questions like ‘Overall, how well did you sleep last night?’ on a scale of 0 – 100 as well as visited a clinic every two weeks to undergo a series of physical sensory tests. The daily diary data (or self-reports) consists of one primary endpoint ‘Overall, how severe have your FM symptoms been today?’ [FM sym] and 13 secondary endpoints: fatigue, sadness, stress, mood, anxiety, satisfaction with life, overall sleep quality, trouble with sleep, ability to think, headaches, average daily pain, highest pain and gastric symptoms (Younger

& Mackey, 2009). Fig. 1 shows data of selected variables for two representative participants. It can be observed that with introduction of drug, the participants report marked changes in pain levels and sleep quality that vary over time. The appropriate description of this dynamical systems response will be the focus of the modeling discussion developed further in the paper.

One of the important issues in data analysis from human subjects, and particularly from clinical trials, is the focus on single subject (idiographic) vs multiple subject (nomothetic) analysis (Moleenaar & Campbell, 2009). From the perspective of adaptive behavioral interventions, the focus in this paper is on performing single subject analysis.

2.2. General description of variables

From an input-output dynamical systems perspective, the variables from the naltrexone trial can be classified as following:

- **Outputs:** There is a clinical interest in understanding the magnitude and speed at which naltrexone affects various FM symptoms during the intervention. Hence typical symptoms like pain, fatigue, sleep disturbance, which correspond to dependent variables in the system, are classified as outputs.
- **Inputs:** Drug and placebo are classified as the primary inputs in this analysis, as they are introduced externally to the system and can be manipulated by the clinician. In addition to these primary inputs, there are other exogenous or disturbance variables affecting the outputs. Variables in the self-reports such as anxiety, stress, and mood are treated as measured disturbance inputs that when coupled with the primary inputs can help better explain the output variance, and ultimately improve the overall goodness-of-fit of the model.

As biological systems are characterized by complex interdependent components, it is difficult to define purely exogenous variables and dependent variables. This interconnection or feedback mechanism (both positive and negative) can result in cross correlation between endpoints and unmeasured noise collected from medical treatments and hence such experiments can be classified, in a classical system identification sense, as closed loop experiments. There may be a relationship between variables such that ‘outputs’ affect ‘inputs’ e.g., an elevated pain condition may affect anxiety levels, although the existence of the feedback path is not clear. In the absence of *a priori* information, this problem is tackled using direct methods by considering it as an open loop system

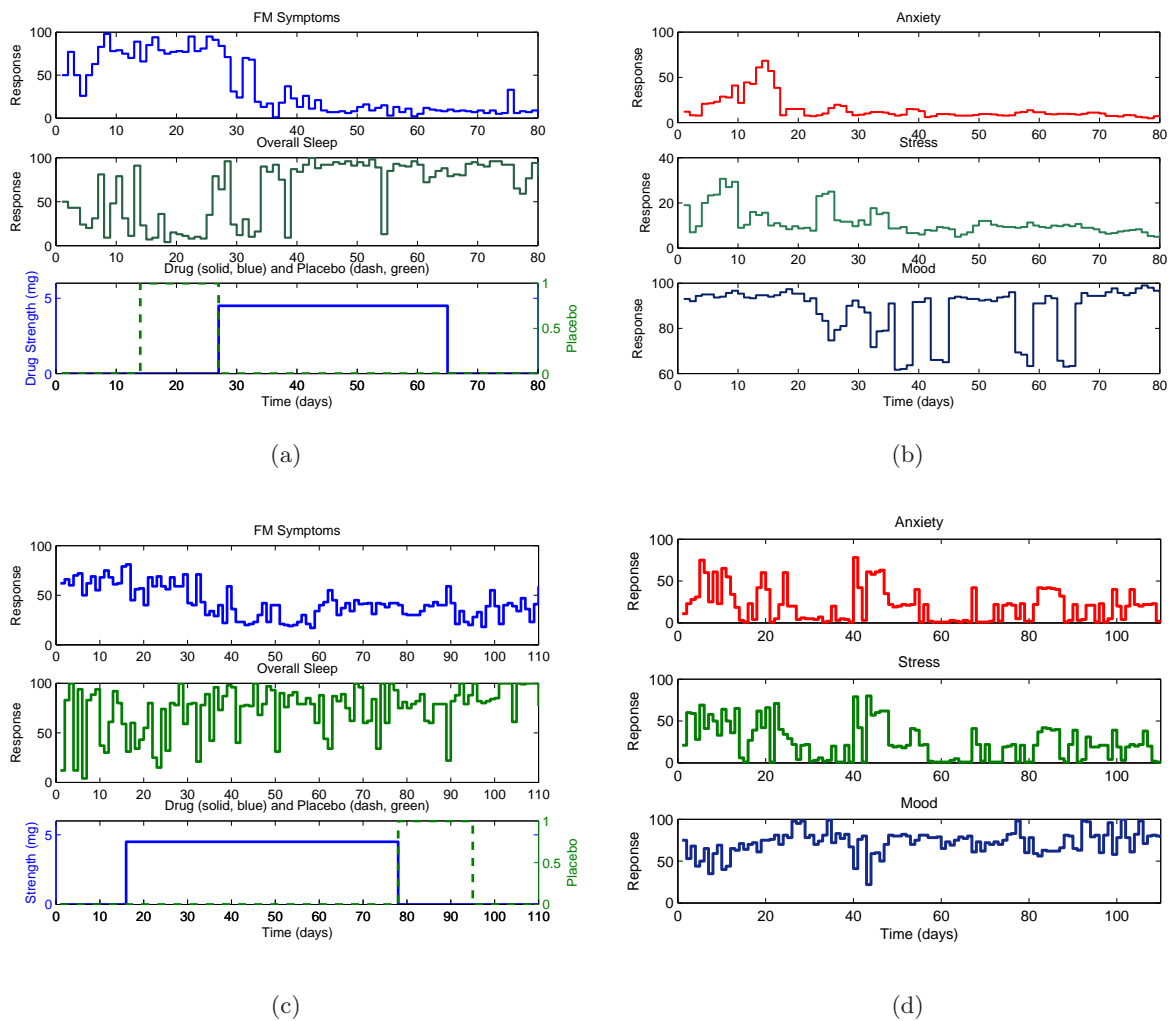


Figure 1: Primary self-report variables associated with naltrexone intervention of fibromyalgia as shown for two representative participants: one participant from the pilot study with placebo-drug (P-D) protocol ((a),(b)) and a participant from the full study with drug-placebo (D-P) protocol ((c),(d)). With the introduction of naltrexone, there is a significant decrease in FM symptoms and increase in sleep quality over time for both participants.

(Ljung, 1999). In the ensuing section, the modeling methodology does not attempt to model the internal mechanisms of FM but rather build an overall response model describing how the drug and external factors affect a number of FM symptoms, so that predictive information can be used by a controller to assign dosages based on measured participant responses.

3. Using system identification to model FM intervention dynamics

In light of the unknown dynamics of FM, an empirical modeling approach is proposed where input-output data of a single participant is used to build a model describing the effect of drug and external factors on FM symptoms.

3.1. System identification procedure

The modeling process undertaken in this study can be summarized in three subparts as follows:

1. *Data preprocessing.* Initially the data is pre-processed for missing entries using a simple mean of immediate neighbors for single missing items, and interpolation for multiple consecutive missing items. To reduce the high frequency content in the time series, a three-day moving average filter $L(q)$ is applied:

$$L(q) = (1 + q^{-1} + q^{-2})/3, \quad (1)$$

where q is the forward shift operator defined as $qy(t) = y(t + 1)$.

2. *Discrete-time modeling using multi-input Autoregressive with exogenous input (ARX) models.* The filtered data is fitted to a parametric multi-input ARX- $[n_a \ n_b \ n_k]$ model:

$$A(q)y(t) = \sum_{i=1}^{n_u} B_i(q)u_i(t - n_{k_i}) + e(t), \quad t = 0, 1, 2, \dots \quad (2)$$

where n_u represents the number of inputs, n_a , n_b and n_k are model orders, $e(t)$ is the prediction error, and $A(q) = 1 + \sum_{j=1}^{n_a} a_j q^{-j}$ and $B_i(q) = \sum_{j=1}^{n_{b_i}} b_j q^{-j+1}$ are polynomials in q . The philosophy is to start with a simpler parametrization (ARX) and add complexity as required. ARX models are computationally simple to estimate and can be consistently estimated provided the inputs are persistently exciting and the model structure is sufficiently high. In examination of multiple participants, ARX-[4 4 1] models were the highest order

needed, and in many cases ARX-[2 2 1] models were suitable (as determined by the classical prediction-error validation criteria, per Ljung (1999); Deshpande (2011)).

The procedure for the choice of input signals is to begin with drug and placebo, which are expected to contribute significantly to FM symptoms for all participants. Additional input variables are then introduced sequentially to improve the goodness of fit. From a statistical perspective, it can be shown that adding extra inputs results in improved covariance of the parameter estimate (stronger for ARX/ARMAX structure) under the assumption that they are independent (Gevers et al., 2006). Consequently, while increasing the number of inputs improves the overall fit, an exceptionally high fit may not necessarily imply a highly predictive model. As the protocol applied in this study did not allow for a crossvalidation data set, proper judgement on the choice of input variables that adequately describes the data across all participants must be made.

3. *Simplification to a continuous time model.* There has been increasing recent interest in identifying continuous time dynamical models directly from sampled data (Garnier & Young, 2014). In this work, the step response from the estimated ARX model serves as the data for estimating a parsimonious continuous second order model structure of the form:

$$G(s) = \frac{K_p(\tau_a s + 1)}{\tau^2 s^2 + 2\zeta\tau s + 1}. \quad (3)$$

From Equation (3), useful system information such as gain, time constant, overshoot, rise and settling times for each input can be obtained which can be used to classify participants as responders or non-responders to the drug. In principle, the order of the continuous time transfer function is problem dependent, but as shown later in the paper, the second order structure according to Equation (3) was sufficient to capture all participants for this dataset. The estimation procedure used corresponds to the `Process Models` routine in MATLAB's System Identification Toolbox (Ljung, 2009; Ljung & Singh, 2012).

The use of prediction-error models, and ARX models in particular, is justified because one can rely on well-established bias relations to obtain insight. For purposes of illustration, consider a system described by one manipulated input (e.g., drug), one measured disturbance input (e.g., anxiety)

and noise, with plant and estimated models as follows:

$$y(t) = p(q)u(t) + p_d(q)d(t) + H(q)\nu(t) \quad (4)$$

$$= \tilde{p}(q)u(t) + \tilde{p}_d(q)d(t) + \tilde{p}_e(q)e(t). \quad (5)$$

The one-step-ahead prediction error can be written as:

$$e(t) = \tilde{p}_e(q)^{-1}(y(t) - (\tilde{p}(q)u(t) + \tilde{p}_d(q)d(t))). \quad (6)$$

Parseval's theorem can be used to relate the filtered prediction-error ($e_F(t) = L(q)e(t)$) with its power spectrum ($\Phi_{e_F}(\omega)$), which is defined as:

$$\Phi_{e_F}(\omega) = \frac{|L(q)|^2}{|\tilde{p}_e(q)|^2} (|p - \tilde{p}|^2\Phi_u(\omega) + |p_d - \tilde{p}_d|^2\Phi_d(\omega) + 2\text{Re}((p - \tilde{p})(p_d - \tilde{p}_d)^*)\Phi_{ud}(\omega) + |H|^2\sigma_\nu^2), \quad (7)$$

where $\nu(t)$, a sequence of independent random variables with zero mean and variance σ_ν^2 , is assumed to be uncorrelated with $u(t)$ and $d(t)$, $L(q)$ is the prefilter, Re denotes the real part of a complex number and $*$ is used to represent its complex conjugate. From Equation (7), it is possible to obtain insights into how input power, model structure, cross-correlation between signals, and other factors can influence the goodness-of-fit in the identification process.

3.2. Case studies

Participant from the pilot study

In this subsection, the focus is on the application of the system identification modeling procedure to a participant from the pilot study, with data as seen in Fig. 1. Equation (7) is used to systematically examine the role of model structure, inputs signal spectral information and input cross correlations on the goodness-of-fit.

- *Model structure.* In this work, ARX models of reasonable dimension were found to be sufficient and no significant improvement was observed with more complex parameterizations.
- *Input cross correlation.* Since various variables are measured in the experiment, the procedure is to choose inputs which have minimum cross spectra ($\Phi_{ud}(\omega)$). The sample cross correlation function (Box et al., 1994) is used to better understand the relationship between different variables and with drug and placebo. For this participant, the headache and gastric variables

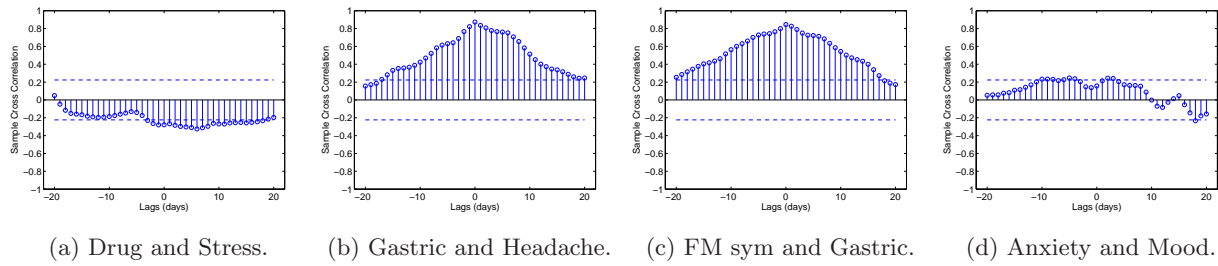


Figure 2: Sample cross correlation function plots between drug and other variables with two standard error bounds over ± 20 lags for participant from the pilot study.

(See Fig. 2b) have a high degree of cross-correlation, and gastric is also correlated with the FM symptoms output (See Fig. 2c). Adding them as inputs did not yield good estimates. In comparison, anxiety and mood (See Fig. 2d) are essentially uncorrelated and offer good estimates when included as inputs. Additional analysis can be found in Deshpande (2011).

The multi-input ARX-[2 2 1] models (with respective input(s) and FM symptoms treated as the primary output) are as follows: 1) **Model 1** (Drug) 2) **Model 2** (Drug, Placebo) 3) **Model 3** (Drug, Placebo, Anxiety) 4) **Model 4** (Drug, Placebo, Anxiety, Stress) 5) **Model 5** (Drug, Placebo, Anxiety, Stress, Mood) 6) **Model 6** (Drug, Placebo, Anxiety, Stress, Mood, Gastric) 7) **Model 7** (Drug, Placebo, Anxiety, Stress, Mood, Gastric, Headache) 8) **Model 8** (Drug, Placebo, Anxiety, Stress, Mood, Gastric, Headache, Life) and 9) **Model 9** (Drug, Placebo, Anxiety, Stress, Mood, Gastric, Headache, Life, Sadness). Fig. 3 shows the corresponding fits for Models 1 through 9, which explain 46.57% to 79.69% of the output variance, respectively. Beyond the five inputs noted in Model 5, one observes that adding more variables does not significantly improve the fit and ultimately (for Model 9) results in overparameterization. Hence, the inputs from Model 5 are used as the base for multi-input ARX models for this participant.

Fig. 4 shows the estimated step responses resulting for the ARX models for the specific case of the naltrexone drug input. The final model has a gain of -2.47 , indicating a nearly 2.5 point drop in the pain report per mg dose of naltrexone. The negative gain for drug classifies this participant as a responder to the treatment. A rise time (T_r) of slightly over 5 days, and a 98% settling time (T_s) of nearly 11.5 days characterizes the naltrexone response for this participant. Table 1 shows how including additional inputs improved the goodness-of-fit for this participant.

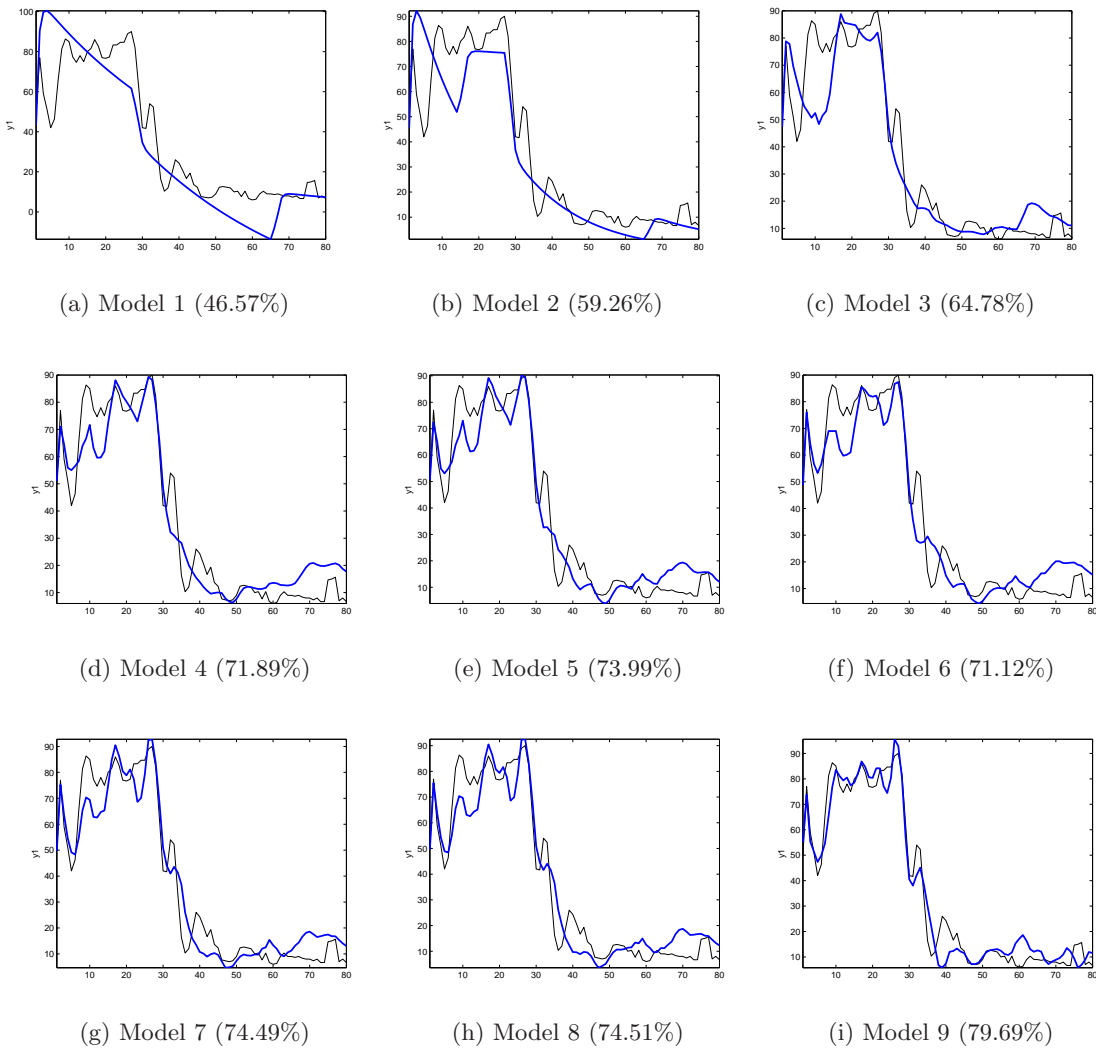


Figure 3: Estimated Models 1 through 9, output vs. actual (FM sym) output using ARX [2 2 1] structure for participant from the pilot study. The percent fits (shown in parenthesis) show no significant improvement beyond Model 5.




Figure 4: ARX model step responses for the drug-FM symptoms transfer function along with the second order continuous time approximation for the case of Model 5. As the model fit improves from Model 1 to 5, the steady-state gain settles at approximately -2.5.

Table 2 summarizes the estimated transfer functions for all five inputs (manipulated and disturbance) of the Model 5 structure. For all these transfer functions, the settling times and rise times (with the exception of Mood-FM) are in close range to each other. The positive gain for the placebo input indicates that in the case of this participant, the administration of placebo has a detrimental effect. The large magnitude of the placebo gain is in part a consequence of binary nature of that input signal. On examining the gains for the measured disturbance models (anxiety, stress, and mood) for this participant, these correspond to 0.86, 2.29, and -0.091 , respectively. The positive values for the anxiety and stress gains and negative for mood agree with the clinical observations that increase in anxiety and stress and decrease in mood should worsen FM symptoms. The low magnitude of the mood gain, coupled with the relatively small contribution of this input to the percent variance described by the model (approximately 2%) indicates the low importance of this variable as a contributor to FM symptoms. Table 2 also includes the model resulting from the effect of drug to overall sleep. The positive gain in this transfer function demonstrates that for this participant, administering naltrexone improved sleep quality. Furthermore, $\tau_a < 0$ for this model denotes the presence of inverse response.

Model	%fit	AIC	K_p, τ, ζ, τ_a	$T_r(\text{days})$	$T_s(\text{days})$
1	46.5	3.64	-12.03, 5.67, 4.14, 21.3	75.5	139.69
2	59.2	3.58	-0.91, 3.5, 2.67, 44.4	0.43	75.06
3	64.7	3.54	-1.02, 2.09, 1.5, 15.3	0.43	25.6
4	71.8	3.42	-3.11, 1.62, 1.24, 0.22	7.53	14.38
5	73.9	3.44	-2.47, 1.57, 1.26, 1.96	5.12	11.49

Table 1: Model estimate summary for the drug-FM model for the pilot study participant. Percent (%) fit and Akaike information criterion (AIC) measure correspond to the multi-input ARX-[2 2 1] model structure.

Model	K_p, τ, ζ, τ_a	$T_r(\text{days})$	$T_s(\text{days})$
Drug-FM	-2.47, 1.57, 1.26, 1.96	5.12	11.49
Placebo-FM	45.81, 1.57, 1.26, 1.15	6.59	13.06
Anxiety-FM	0.86, 1.57, 1.26, 0.24	7.45	14.24
Stress-FM	2.29, 1.57, 1.26, 0.49	7.31	13.94
Mood-FM	-0.091, 1.57, 1.26, 4.67	0.8	11.93
Drug-Overall Sleep	4.98, 2.13, 1.04, -3.35	7.06	15.83

Table 2: Model parameter tabulation for various inputs-FM continuous models as well as the drug-overall sleep (Drug-Overall Sleep) model for pilot study participant. The participant shows reduction in pain and improvement in sleep with drug intake.

Full study participants

Table 3 shows a summary of participant responses from the full study for Placebo-Drug (P-D) protocol. In general for participants from the full study, additional inputs (such as sadness and headache) as well as ARX models with higher orders (e.g., [4 4 1]) were required for improved fits. The case for most participants was that inputs corresponding to Model 9 gave a better fit. For a total 15 participants in this protocol, nine were classified as responders and six as non-responders based on the estimated model gain for drug as input and FM sym as output ($K_p(\text{DFM})$). Settling time (in days) for each case are also noted. The gains for each participant are shown with one standard deviation on estimated gains (from system identification) whereas, in rows corresponding

to average values, the deviation of mean of each participant's response is noted. It can be noted that the average model gain ($K_p(\text{DFM})$) for responders was -3.56 whereas for non-responders was 2.06 . Overall, for all participants of this protocol the gain was -1.31 . Similarly gain for drug as input and overall sleep as output ($K_p(\text{DS})$) together with the settling time was also tabulated and it was noted that the response of sleep to drug was not strong in many cases with average gain for all participants being -0.23 . In general, the drug response of full study participants was weaker compared to cases in the pilot study as can be noted by the gain magnitudes and large error bounds which make the classification of participants difficult. Details of participant responses for the drug-placebo (D-P) protocol can be found in Deshpande (2011).

3.3. Model validation

The following standard methods are used to validate the estimated models (Ljung, 1999):

- *Residual analysis.* A residual analysis is conducted on all estimated models using auto correlation of the residual and cross correlation between the inputs and residual. For the majority of the participants in this study, ARX-[2 2 1] or [4 4 1] models met the classical prediction error criteria.
- *Step responses from estimated ARX models.* After a model has passed residual analysis, the model step responses are analyzed. From Model 2 onwards in Fig. 4, the responses tend to settle rapidly to a steady state with improved goodness-of-fit.

While it would have been desirable to have applied cross-validation to this analysis, when the data was partitioned into estimation and validation sets, it was found to lack the excitation required to support multi-input crossvalidation. This is a consequence of the limited number of data points in this study and the experimental procedure that was followed for the naltrexone and placebo dosages. Instead, the percent improvement with additional inputs added to the model are noted, while simultaneously examining the Akaike information criterion (AIC) values and step response results to avoid the consequences of overparameterization (e.g., large changes in gain and settling time with parameters that do not agree with physical insight).

In summary, the majority of the pilot study participants were adequately modeled with Model 5 using the ARX-[2 2 1] structure and in the full study with Model 9 using the ARX-[4 4 1] structure, and subsequent approximation to the second order continuous time model. It was observed that

Responders						
#	% fit(DFM)	K_p (DFM)	T_s (DFM)	% fit(DS)	K_p (DS)	T_s (DS)
1	64.92*	-4.70*±7.67	23.23*	72.50	-3.28±6.98	27.24
2	48.71	-0.66±1.15	15.28	35.19	0.668±1.52	16.12
3	29.82	-0.6±1.71	15.89	23.17	-2.83±4	45.55
4	29.23#	-11.66#± 12.36	67.01#	28.89	1.59±2.7	13.02
5	18.08	-0.83±2.3	10.51	29.87	-4.89±3.7	20.39
6	54.70	-2.79±1.88	23.11	17.93	0.4±2.77	11.57
7	40.00	-2.27±3.99	27.1	43.38	-0.92±1.16	17.19
8	54.13	-8.44±5.12	12.63	64.55	-1.88±1.46	19.11
9	59.00	-0.112±0.83	16.04	44.58	2.27±0.77	13.05
Average Values (std. deviation for nominal gain only)						
—	44.28±15.8	-3.56±4	23.42±17.2	40±18.4	-0.98±2.4	20.36±10.5
Non Responders						
10	44.68	3.96±3	33.05	—	—	—
11	2.19*	4.09*±3.96	39.21*	20.57	-2.53±2.47	15.54
12	27.99	0.82±1.73	30.37	38.45	1±1.16	17.8
13	31.17	1.94±2.07	34.97	22.47	6.23±4.94	42.46
14	40.31	0.37±1.24	16.24	26.7	0.85±1.33	22.67
15	34.57	1.17±2.85	15.4	49.14	0±1.16	18.34
Average Values (std. deviation for nominal gain only)						
—	30.15±14.9	2.06±1.6	28.2±10	31.46±12	1.11±3.17	23.62±10.9
Total Average Values (std. deviation for nominal gain only)						
—	38.63±16.5	-1.31±4.27	25.33±14.5	36.95±16.4	-0.23±2.77	21.43±10.4

Table 3: Tabulation of system responses to drug for the placebo-drug (P-D) protocol for selected participants from the full study. Corresponding model fits are also noted. Model 9 was used unless noted otherwise as * implying input combination as {Drug, Placebo, Anxiety, Stress, Mood, Life, Sadness and Gastric} and # where inputs corresponding to Model 7 are used. In cases of bad data (participant 10), no model was estimated.

the response to drug was stronger in the pilot study as compared to the full study; likewise, a large number of full study participants showed placebo response. For those participants whose modeling results would classify them as responders to treatment, the estimated models are used as the basis for adapting the intervention using a control engineering approach.

4. Model predictive control of naltrexone intervention for fibromyalgia

In this work, model predictive control (MPC) is used as the algorithmic framework for making systematic dosage assignments. This control technology effectively combines the feedback-feedforward control action by on-line optimization of a cost function using a receding horizon (Qin & Badgwell, 2003). The MPC approach allows for flexibility to integrate critical clinical objectives and constraints, and hence is particularly suited for designing treatment regimens for adaptive interventions. It has seen medical applications from diabetes mellitus control to HIV/AIDS management (Zurakowski & Teel, 2006; Wang et al., 2010). To achieve a desired performance, a three-degree-of-freedom (3 DoF) approach is used to tune the controller (Lee & Yu, 1994; Wang & Rivera, 2008). This tuning methodology enables performance requirements associated with setpoint tracking, anticipated measured disturbance rejection and unmeasured disturbance rejection to be adjusted independently by varying parameters α_r , α_d and f_a respectively (as discussed later). These parameters can be adjusted between values 0 and 1; they in turn specify the response of a filter which supplies a filtered signal to the controller (α_r for setpoint tracking and α_d for measured disturbance rejection) or adjust the observer gain (K_f) for unmeasured disturbance rejection (f_a). The hybrid extension of this 3 DoF approach (Nandola & Rivera, 2013) is considered here.

4.1. Clinical goals

Adaptive interventions employ decision rules and repeated assessments of participant response to improve outcomes (Collins et al., 2004). In a control engineering approach to adaptive interventions (Rivera et al., 2007; Zafra-Cabeza et al., 2011; Nandola & Rivera, 2013), the controller assigns dosages to each participant as dictated by model dynamics, problem constraints, and disturbances (both measured and unmeasured). The control system aims at functionally performing the following three tasks:

- *Setpoint tracking.* Drug dosages are assigned to take an outcome of interest (such as FM symptoms or overall sleep quality) to a desired goal or setpoint. For example, a clinician may decide on a goal of 45% reduction in general pain symptoms within two weeks of drug treatment.
- *Measured disturbance rejection.* The controller adjusts drug dosages to mitigate the effect from *reported* external influences using estimated disturbance models. For instance, if some

external event (e.g., reported anxiety) that leads to elevated FM symptoms is known *a priori*, then dosages can be adjusted in anticipation to compensate for that disturbance.

- *Unmeasured disturbance rejection.* The controller adjusts drug dosages to mitigate the effect of unknown and unmodeled external influences. For example, a sudden unexpected event may result in worsening pain (FM symptoms). In such cases, the controller alters dosages to mitigate the unmeasured disturbance.

In addition to accomplishing the three functional modes of the control system, a number of practical clinical requirements have to be integrated into the controller design. These are:

- *Limits on possible dosages.* Drug efficacy is generally defined on some dosage levels, as well as dosage magnitudes has to be limited to certain bounds to ensure safe usage. In this work, the drug naltrexone is varied for simulation purposes as

$$0 \leq u(k) \leq 13.5 \text{ mg}, \quad k \in \{0, 1, \dots, N - 1\}, \quad (8)$$

where N is the length of simulation.

- *Gradual change in dosages.* Dosage changes should not be very abrupt due to concerns with drug withdrawal and toxicity. This translates into a limit on the move size of the signal as:

$$\Delta u_{\min} \leq |u(k) - u(k - 1)| \leq \Delta u_{\max}, \quad (9)$$

- *Categorical or discrete dosages.* This clinical consideration necessitates the use of a hybrid approach. In this work, the treatment is confined to eight drug dosage levels defined as $u(k) \in \mathbb{I} = \{0, 1.92, 3.85, 5.76, 7.68, 9.6, 11.58, 13.5\}$ mg. Subsequently in this paper, the term *hybrid* MPC is used to define a controller which assigns dosages categorically, whereas a *continuous* MPC will assign dosages anywhere in the given range.

4.2. MPC problem formulation

The details of the proposed MPC formulation are now presented. As was pointed out earlier, an important consideration in adaptive interventions is that intervention dosages can assume only discrete levels, and therefore it is necessary to consider hybrid algorithms (Nandola & Rivera, 2013). A mixed logical dynamical (MLD) framework is used to represent linear hybrid systems

which are systems with real and integer states, inputs and constraints (Bemporad & Morari, 1999), as shown:

$$x(k+1) = Ax(k) + B_1u(k) + B_2\delta(k) + B_3z(k) + B_d d(k) \quad (10)$$

$$y(k) = Cx(k) + d'(k) + \nu_n(k) \quad (11)$$

$$E_2\delta(k) + E_3z(k) \leq E_5 + E_4y(k) + E_1u(k) - E_d d(k), \quad (12)$$

where, in general, $x \in \mathbb{R}^{n_x}$ and $u \in \mathbb{R}^{n_u}$ represent states and inputs of the system. $y \in \mathbb{R}^{n_y}$ is the output and $d \in \mathbb{R}^{n_d}$, $d' \in \mathbb{R}^{n'_d}$ and $\nu_n \in \mathbb{R}^{n_\nu}$ represent measured disturbances, unmeasured disturbances and measurement noise signals respectively. $\delta \in \{0, 1\}^{n_\delta}$ and $z \in \mathbb{R}^{n_z}$ are discrete and continuous auxiliary variables respectively, which along with the input u , output y and disturbance d form the linear inequality constraint shown in Equation (12) in order to enforce logical/discrete decisions. The effect of all unmeasured disturbances is lumped as d' in the measurement equation.

A standard quadratic cost function ($\|\star\|_{Q_\star}^2 = (\star)^T Q_\star (\star)$) is used to calculate the decision vector as:

$$\begin{aligned} \min_{\{[u(k+i)]_{i=0}^{m-1}, [\delta(k+i)]_{i=0}^{p-1}, [z(k+i)]_{i=0}^{p-1}\}} J \triangleq & \sum_{i=1}^p \|(y(k+i) - y_r)\|_{Q_y}^2 + \sum_{i=0}^{m-1} \|(\Delta u(k+i))\|_{Q_{\Delta u}}^2 \\ & + \sum_{i=0}^{m-1} \|(u(k+i) - u_r)\|_{Q_u}^2 + \sum_{i=0}^{p-1} \|(\delta(k+i) - \delta_r)\|_{Q_d}^2 + \sum_{i=0}^{p-1} \|(z(k+i) - z_r)\|_{Q_z}^2 \end{aligned} \quad (13)$$

$$\text{s.t. } y_{\min} \leq y(k+i) \leq y_{\max}, \quad i \in \{1, \dots, p\} \quad (14)$$

$$u_{\min} \leq u(k+i) \leq u_{\max}, \quad i \in \{0, \dots, m-1\} \quad (15)$$

$$\Delta u_{\min} \leq \Delta u(k+i) \leq \Delta u_{\max}, \quad i \in \{0, \dots, m-1\}, \quad (16)$$

and also subjected to state and output equations with mixed integer constraints as shown in Equations (10)-(12), where p is the prediction horizon and m is the control horizon. The vector 2-norm are weighted by matrices Q_\star as in Q_y , $Q_{\Delta u}$, Q_u , Q_d , and Q_z are the penalty weights on the error, move size, control signal, auxiliary binary variables and auxiliary continuous variables, respectively. The problem is formulated as a tracking control system using references y_r , u_r , δ_r and z_r for output, input, discrete and continuous auxiliary variables, respectively.

4.3. Controller tuning

The tuning approach for HMPC can be discussed broadly in two sub topics:

- *Setpoint tracking and measured disturbance rejection.* In the 3 DoF approach, the speed and shape of setpoint tracking and measured disturbance rejection can be directly manipulated using filters, and hence the closed loop response can be varied to achieve a desired performance. The choice of the filter depends on whether the set point or disturbances changes are asymptotically step (Type I) or ramp (Type II). In this work, a Type I filter is used for both cases, described as:

$$f(q, \alpha_{r,d}^j) = \frac{(1 - \alpha_{r,d}^j)q}{q - \alpha_{r,d}^j}, \quad j \in \{1, \dots, n_y\}, \quad (17)$$

where $\alpha_{r,d} \in [0, 1)$ and n_y is the number of outputs. The tuning parameter to alter the setpoint response is denoted by α_r , while the tuning parameter for measured disturbance rejection is α_d . Hence the controller can be tuned for slower rejection of measured disturbances, for example, by more extensive filtering of the disturbance signals.

- *State estimation (unmeasured disturbance rejection).* Considering the stochastic and uncertain nature of self-reports and estimated models, the control system is augmented with a state observer to mitigate the errors associated with system prediction. In this system, unmeasured disturbances can occur externally or can originate from plant-model mismatch. A parametrized observer (Lee & Yu, 1994; Wang & Rivera, 2008; Nandola & Rivera, 2013) is used, where the observer gain can be written as:

$$K_f = \begin{bmatrix} 0 & F_b^T & F_a^T \end{bmatrix}^T \quad (18)$$

$$F_a = \text{diag}\{(f_a)_1, \dots, (f_a)_{n_y}\} \quad F_b = \text{diag}\{(f_b)_1, \dots, (f_b)_{n_y}\} \quad (19)$$

$$(f_b)_j = \frac{(f_a)_j^2}{1 + \bar{\alpha}_j - \bar{\alpha}_j(f_a)_j}, \quad j \in \{1, \dots, n_y\}, \quad (20)$$

and where $(f_a)_j$ is a tuning parameter that lies between 0 and 1. In this work, as asymptotically step inputs are considered, hence $\bar{\alpha}_j$ is equal to zero (Wang & Rivera, 2008). As $(f_a)_j$ approaches zero, the state estimator increasingly ignores the prediction error. In contrast, as $(f_a)_j$ approaches 1, the state estimator tries to compensate for all prediction error and hence may cause the controller response to become extremely aggressive.

Full details of the controller formulation, comparison between different tuning settings, and applications to a preventive behavioral intervention and inventory management in supply chains can be found in Nandola & Rivera (2013).

5. Closed-loop simulation results

This section demonstrates closed-loop drug dosage assignment on the representative participant from the pilot study discussed in Section 3.2. The system under consideration is a five input, single output dynamical model with one manipulated variable (naltrexone), four disturbance variables (placebo, anxiety, stress and mood) and one output (FM symptoms). FM symptoms serves as the primary outcome in the analysis ($n_y = 1$), while anxiety (assumed to be reported daily by the participant) serves as the measured disturbance signal. All other disturbances are set equal to zero with no measurement noise ($\nu_n(k) = 0$). The drug dosages range from 0 to 13.5 mg with eight possible values. The eight discrete inputs form an arithmetic progression and can be represented logically in the MLD framework as:

$$\delta_i(k) = 1 \Leftrightarrow z_i(k) = 13.5 - (i - 1) \times 1.9286 \quad i \in \{1, 2, \dots, 8\} \quad (21)$$

$$u(k) = \sum_{i=1}^8 z_i(k), \quad \sum_{i=1}^8 \delta_i(k) = 1. \quad (22)$$

These conditions and implications (\Leftrightarrow) are then converted into inequality constraints as represented in Equation (12). The simulation parameter values are as follows: the prediction horizon (p) is set to 25 days (chosen larger than the open loop settling time from the estimated Model 5 (see Table 1)) and the control horizon (m) to 15 days (chosen to be long enough for a feasible solution but less than the prediction horizon). As the 3 DoF approach is used to tune the controller, the objective weights will not be adjusted and are fixed as follows: $Q_y = 1$, $Q_{\Delta u} = Q_u = Q_d = Q_z = 0$. The three “knobs” in the controller are varied as: $\alpha_r \in [0, 1)$, $\alpha_d \in [0, 1)$, and $f_a \in (0, 1]$. These simulation parameters are now used to generate closed-loop control results which are shown in two sections:

1. Nominal performance for tracking and disturbance rejection (Section 5.1),
2. Robust performance under plant-model mismatch (Section 5.2).

In this hypothetical simulation of a control-oriented naltrexone adaptive intervention, the goal is to demonstrate how the 3 DoF formulation gives the flexibility to achieve a broad range of time-domain responses, and hence the controller can be tuned as per clinical requirements. An initial general proof-of-concept is followed by an analysis involving more rigorous performance metrics.




Figure 5: Comparison of closed-loop responses obtained from an eight level hybrid MPC (solid, blue) with continuous MPC (dash, black) when tuning parameters are $(\alpha_r, \alpha_d, f_a) = (0.5, 0.5, 0.5)$. The setpoint is shown by a dashed red line.

5.1. Nominal performance

This section evaluates the performance of the controller when a true model of the system is known. Three independent events take place in the simulation: setpoint tracking starts at $k = 0$, a measured disturbance acts at $20 \leq k \leq 40$ with magnitude 16.52 and an unmeasured disturbance at $k = 55$ of magnitude 9.63. The output variable starts with a baseline value of 50 and a change of -9.5 is applied at $k = 0$ as shown by the reference in Fig. 5. The result is shown for tuning parameters $(\alpha_r, \alpha_d, f_a) = (0.5, 0.5, 0.5)$. The length of the simulation is $N = 75$ days.

MPC where $u(k)$ can take any continuous value on its range and with no filtering represents the best possible performance by the controller. Clinically, this can be an initial benchmark which can be used to get a sense for the treatment regimen obtained from the control system. Next, depending on the constraints on drug dosage levels, a new treatment regimen has to be generated which can be contrasted with the continuous case. However, the drug dosage changes may be perceived as too aggressive by the clinician and hence, based upon the exact requirements (e.g., pain reduction by 30%), the hybrid controller can be de-tuned using 3 DoF tuning variables. For setpoint tracking, α_r can be adjusted to suit the expected response. Similarly, the response to

disturbances can be varied by α_d and f_a to suit the conditions at hand. In general, by increasing the filtering action, dosage changes are smoother and more clinically acceptable. For measured disturbances, continuous MPC can offer better compensation than hybrid MPC as the later may be less effective due to categorical constraints. At $k = 55$ an abrupt change in the pain report occurs due to an unmeasured disturbance; this change is not part of the model prediction. The controller reacts by increasing the drug dosage to compensate.

In practice, setpoint tracking has to be accomplished in presence of disturbances and hence it is now shown that the control algorithm used here performs well in events occurring simultaneously. Both measured and unmeasured disturbances (of magnitude as before) act at $k = 0$ as shown in Fig. 6. The length of the simulation is $N = 20$ days. The tracking is considered under both measured and unmeasured disturbances, where a smoother (but more aggressive controls) response is obtained for $\alpha_d = 0, f_a = 1$ whereas the less aggressive controls for tuning $\alpha_d = 0.9, f_a = 0.2$. In both of the cases, the variation of α_r from 0 to 0.8 (as seen on the Y-axis) results in more sluggish setpoint tracking speed for fixed set of tuning parameters α_d, f_a . It can also be observed that to compensate the effect of disturbance, higher dosages are required. In order to quantify the achieved performance under different sets of tunings, two metrics are defined related to the square of 2–norm of the tracking error e_c and the square of 2–norm of change in control Δu as:

$$J_e = \sum_{k=0}^{N-1} e_c(k)^T e_c(k), \quad J_{\Delta u} = \sum_{k=1}^{N-1} \Delta u(k)^T \Delta u(k) \quad (23)$$

where $e_c(k) = y(k) - y_r$ and $\Delta u(k) = u(k) - u(k-1)$. These metrics are tabulated in Table 4, where it can be observed that de-tuning the controller (i.e. decreasing f_a and increasing α_r, α_d) results in higher J_e (more error) but lowering of $J_{\Delta u}$ (less aggressive controls).

So far, only deterministic setpoint and disturbance signals have been considered. A scenario of significant practical interest is when the unmeasured disturbance is of stochastic nature; this is now introduced through an AutoRegressive Moving Average (ARMA) model to evaluate the controller performance. By changing the observer gain through f_a , it is possible to influence disturbance rejection. Consider an ARMA model as follows:

$$d(k) = \frac{(q - 0.3)}{(q - 0.7)^2} a(k) \quad \forall k, \quad a(k) \sim \mathcal{N}(0, 30), \quad (24)$$

where d and a are discrete-time signals with q as the forward shift operator. A realization of this ARMA noise is shown in Fig. 7a where the unmeasured disturbance varies on a scale of 100. The

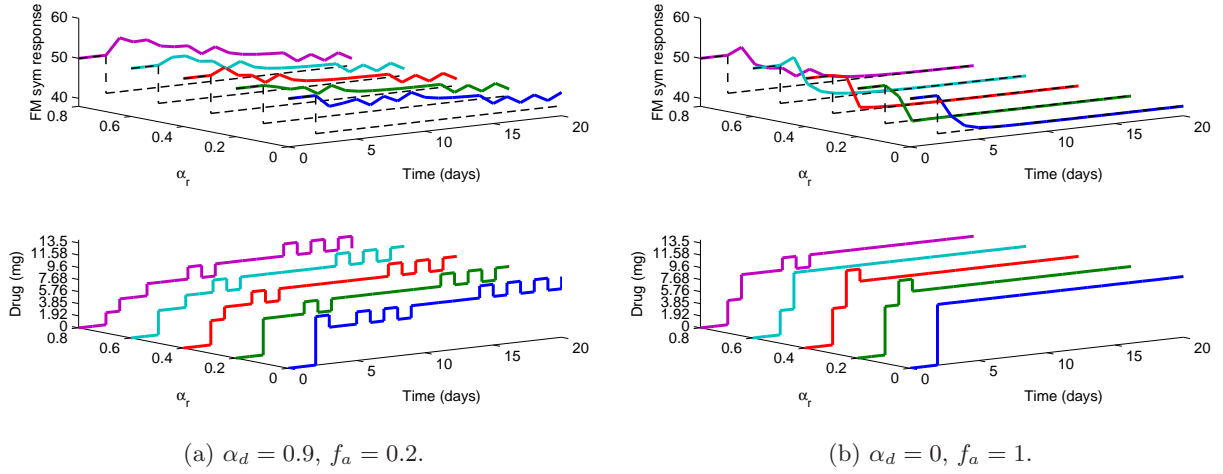



Figure 6: Comparison of setpoint (dotted black lines) tracking under both unmeasured and measured disturbance for various tuning (α_r) values under fixed α_d and f_a .

$\alpha_d = 0.9, f_a = 0.2$			$\alpha_d = 0, f_a = 1$		
α_r	J_e	$J_{\Delta u}$	α_r	J_e	$J_{\Delta u}$
0	602.93	107.86	0	108.41	92.98
0.2	628.82	63.22	0.2	131.51	78.1
0.4	710.67	48.35	0.4	167.88	70.66
0.6	855.79	48.35	0.6	235.03	48.35
0.8	1181.53	44.63	0.8	327.85	40.91

Table 4: Performance indices for the e_c and Δu signals under different tuning values for setpoint tracking under disturbances. It can be observed that increases in α_r (implying more filtering action) lead to less aggressive control responses.

length of the simulation is $N = 90$ days.

The two simulation cases are shown in Figure 7b. Two settings are considered to change the observer gain: $f_a = 1$ (which is more aggressive) and then $f_a = 0.1$ (which is more sluggish). Table 5 compares the performance indices related to the error and change in control from both simulations. It can be noted that under $f_a = 1$, a lower error performance (J_e) is observed but results in a trade off with more aggressive control action, whereas using a lower tuning value of $f_a = 0.1$ results in higher J_e , with the corresponding less change in control ($J_{\Delta u}$). To conclude,



(a) Time series of ARMA (2,1) noise per Equation (24). (b) Stochastic unmeasured disturbance rejection.

Figure 7: (a) Time series realization of unmeasured disturbance from the ARMA noise model. (b) Performance of hybrid MPC under different tunings: $f_a = 1$ (dash, black) and $f_a = 0.1$ (solid, blue) when the unmeasured disturbance is a realization of ARMA noise model. The setpoint is shown by a dashed red line.

$f_a = 0.1$		$f_a = 1$	
J_e	$J_{\Delta u}$	J_e	$J_{\Delta u}$
9211.26	55.79	6203.99	1205.08

Table 5: Performance index of the hybrid MPC under stochastic unmeasured disturbance. The observer gain can be varied to obtain a trade off between the metric of tracking error (J_e) and control moves ($J_{\Delta u}$).

the tuning approach gives the user enough flexibility to choose the speed of response, for a given set of weight matrices, to satisfy desired clinical requirements.

5.2. Robust performance

It was previously noted the importance of having a suitable controller under conditions of model uncertainty and participant variability. That scenario is now addressed where the performance of the controller is evaluated when an erroneous model is obtained from the identification procedure. The approach used in this paper is to showcase robustness via simulations using different model uncertainties for a fixed controller tuning. As before, one input-output (drug-FM) and one disturbance (anxiety-FM) model are considered. To simulate model uncertainties, different parametric uncertainties on the estimated models (as shown by Equation (3)) are used to generate scenarios for plant-model mismatch. The plant and disturbance model variations are chosen to illustrate a rich

set of dynamic responses and variability expected from a cohort of participants. For an individual treatment case, these uncertainties may correspond to unmodeled dynamics. These cases are noted in Table 6 for the drug-FM model and anxiety-FM models. The modeling errors in the plant will be compensated through feedback action alone, while modeling errors in the disturbance model will be partially compensated by anticipated feedforward action; the disturbance not compensated by the anticipation will enter the feedback loop as an unmeasured disturbance. Due to constraints on the input drug dosage levels, the uncertainties have not been chosen to be arbitrarily very large.

The observer gain is adjusted through f_a and other parameters are kept constant ($\alpha_r = \alpha_d = 0$). The setpoint is kept constant at 50 and a measured disturbance is applied at $k = 2$ of magnitude 11.05. The length of the simulation is $N = 35$ days. It is important to mention that these results are displayed for clinical inferences using two functional groupings:

1. *When a fixed nominal model is used.* A nominal model is used as a basis by the controller to assign dosages for different plants (the traditional robustness scenario). Clinically, this can be interpreted in two ways: first, that the (estimated) nominal model is an approximation of the true system (and hence the different scenarios represent different uncertain plants) for an individual participant and second, the nominal model represents an average or representative model for a population of participants (and hence the different scenarios represent different participants). This is shown in Fig. 8a-8b.
2. *When the true plant serves as the nominal model.* For each scenario considered in the previous case, the true plant is supplied as the nominal model to the controller. This case can be understood as when accurate modeling (through system identification or otherwise) has been performed for each individual in a population. The key motivation for this is to get a clinical insight as the user can now compare how the controls will vary under plant-model mismatch and on the same page, it will help the user to assess the case when the correct model is available to the controller. Hence, the user can gauge the resultant change in dosing strategies due to modeling errors. This is shown in Fig. 8c.

In Fig. 8 each scenario (on ‘Y’ axis) represents five cases of uncertainty combinations of perturbations in both plant and disturbance model respectively. The following list shows the uncertainty cases $[p, p_d]$ used for simulation from Table 6: a) Scenario 1: [1,1] b) Scenario 2: [4,5] c) Scenario 3: [3,3] d) Scenario 4: [5,3] e) Scenario 5: [4,3].

Case (p)	ΔK_p	$\Delta \zeta_p$	$\Delta \tau_p$	$\Delta(\tau_a)_p$	Case (p_d)	ΔK_d	$\Delta \zeta_d$	$\Delta \tau_d$	$\Delta(\tau_a)_d$
1	0	0	0	0	1	0	0	0	0
2	-29.1%	0	0	-23.46%	2	-15.1%	0	0	66.66%
3	29.1%	0	0	-23.46%	3	15.1%	0	0	66.66%
4	-29.1%	16.6%	259%	-23.46%	4	-15.1%	16.6%	259%	66.66%
5	29.1%	79.3%	-29.1%	-23.46%	5	15.1%	79.3%	191%	66.66%

Table 6: Tabulation of parametric perturbations for the drug-FM (plant p) and anxiety-FM (disturbance p_d) models.

As can be noted, Scenario 1 in all plots is the case of nominal model (and hence no steady state error). When f_a is changed from $f_a = 0.2$ (See Fig. 8a) to $f_a = 1$ (See Fig. 8b), more aggressive control are obtained, although this also results in less output overshoot (max. pain) as noted in Table 7. In Fig. 8c, all the cases are compared with respective scenarios where a correct nominal model is available and it can be observed that the control is better as expected. It can be noted that only one plot is shown for correct nominal model case as both tuning values of f_a result in the same response (no prediction error). Table 7 records the performance indices of this analysis; for $f_a = 0.2$, J_e is higher when compared case-to-case with performance under tuning $f_a = 1$. However, using larger values of f_a may result in aggressive control as in noted by some values of $J_{\Delta u}$. The metrics when a correct nominal model is available are also noted where, as expected, J_e is lower with corresponding large values of $J_{\Delta u}$.

6. Summary and conclusions

This paper demonstrates the design of an adaptive intervention that relies on system identification modeling and hybrid model predictive control to assign appropriate dosage levels of naltrexone as a treatment for fibromyalgia, a chronic pain condition. The approach described in this work generates models from clinical data and assigns categorical dosages by considering hybrid dynamics in a mixed logical dynamical (MLD) framework.

Given the absence of first principles models, a secondary data analysis was performed to estimate parsimonious models from data available through clinical trials. Low-order multi-input ARX models were estimated and approximated to continuous-time second order models. The effect of


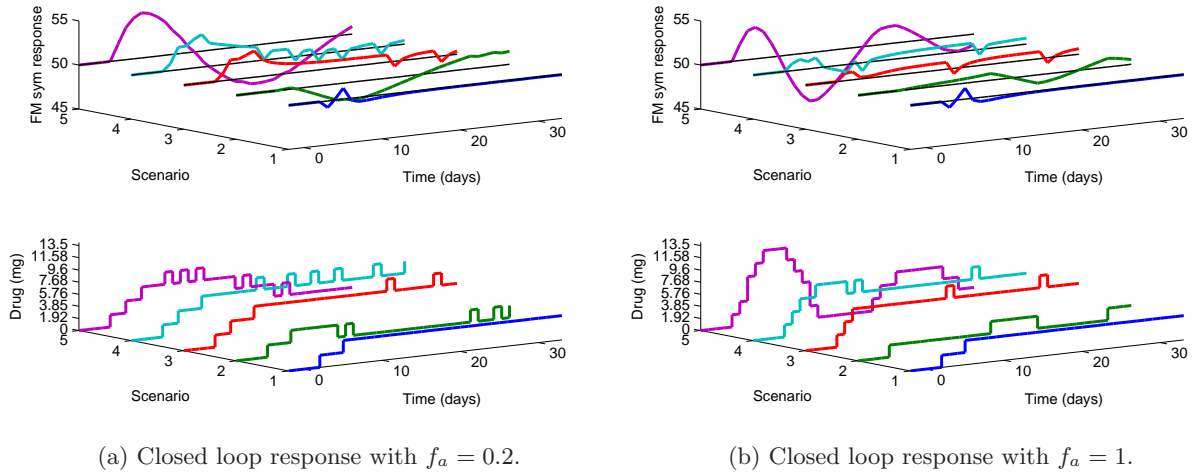


Figure 8: Robustness evaluation when both plant model (drug-FM) and disturbance model (anxiety-FM) are perturbed under measured disturbance, with tuning $\alpha_r = \alpha_d = 0$ under different f_a . Plots (a) and (b) show the closed loop response with under model mismatch ($p \neq \tilde{p}$, $p_d \neq \tilde{p}_d$ where Scenario 1 represents the nominal model) under different tuning. Plot (c) shows the response under no plant-model mismatch.

Tuning	Uncertainty Case (plant (p), disturbance(p _d))	$p \neq \tilde{p}, p_d \neq \tilde{p}_d$			$p = \tilde{p}, p_d = \tilde{p}_d$		
		J_e	$J_{\Delta u}$	Max. pain	J_e	$J_{\Delta u}$	Max. pain
$\alpha_r = \alpha_d = 0, f_a = 0.2$	1,1	2.88	7.43	51.25			
	4,5	44.47	37.19	51.55			
	3,3	48.05	26.03	52.99			
	5,3	75.48	52.07	53.7	2.88	7.43	51.25
	4,3	307.28	55.79	55.12	0.39	126.46	50.23
$\alpha_r = \alpha_d = 0, f_a = 1$	1,1	2.88	7.43	51.25	6.78	63.23	50.56
	4,5	9.05	14.87	50.97	8.55	70.66	50.72
	3,3	12.04	26.03	50.91	0.25	252.92	50.17
	5,3	14.44	33.47	51.17			
	4,3	216.33	78.1	53.57			

Table 7: Performance indices under two sets of tuning for robust performance with different scenarios of plant and disturbance model perturbations.

drug, placebo and other variables on outcomes of interest such as general FM symptoms were systematically included in the modeling procedure. The model yields the dynamical information that can be used to classify participants as responders or non responders to treatment, and to make dosage changes over time.

Models from a representative participant in the clinical trial were used to show how model predictive control can be applied to assign dosages in the presence of disturbances and model uncertainties. The use of an improved 3 DoF tuning approach was demonstrated to give flexible independent tuning for a desired controller performance. The control results were broadly classified under nominal performance and robust performance. Under nominal performance, it is shown how by varying three tuning parameters (α_r, α_d, f_a) related to filters and state observer dynamics, independent tuning for setpoint tracking, measured and unmeasured (deterministic and stochastic) disturbance rejection can be achieved, and their relationship to clinical goals in the intervention. In

the robust performance evaluation, the model parameters are perturbed to create conditions of a plant-model mismatch which would be indicative of participant variability during the intervention.

The results presented in this paper can impact not only the treatment of fibromyalgia, but also the treatment of other chronic pain conditions and the development of adaptive behavioral interventions in general (Dong et al., 2014; Timms et al., 2014). While the focus of this paper was on black-box modeling, an increasing interest in using behavior change theories to inform “just-in-time” decisions in an intervention is receiving increasing activity (Rivera & Jimison, 2013; Martin et al., 2014). It is envisioned that control engineering concepts will play a crucial role in novel individualized treatments, where a closed-loop system can adjust treatment dosages based on daily patient reports of pain and other symptoms of importance. In this way, adaptive interventions relying on control systems engineering can be seen as a cost-effective and efficient method for accomplishing personalized pain interventions. In conclusion, as models are obtained from system identification methods, the effective design of experiments for generating informative clinical data becomes increasingly critical for success of these interventions (Deshpande et al., 2012). A discussion of how optimal input design can play a role in the development of individualized trials while considering clinical requirements is discussed in Deshpande (2014).

7. Acknowledgments

Support for this work has been provided by the Office of Behavioral and Social Sciences Research (OBSSR) of the National Institutes of Health (NIH) and the National Institute on Drug Abuse (NIDA) through grants R21 DA024266 and K25 DA021173. The content is solely the responsibility of the authors and does not necessarily represent the official views of OBSSR, NIDA, or the NIH. Jarred Younger received support from the American Fibromyalgia Syndrome Association (AFSA). Insights provided by Linda M. Collins and Jessica Trail of the Methodology Center, Penn State University during the conduct of this research are greatly appreciated.

References

- Bemporad, A., & Morari, M. (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica*, *35*, 407–427.
- Boissevain, M. D., & McCain, G. A. (1991a). Toward an integrated understanding of fibromyalgia syndrome. I. Medical and pathophysiological aspects. *Pain*, *45*, 227 – 238.

- Boissevain, M. D., & McCain, G. A. (1991b). Toward an integrated understanding of fibromyalgia syndrome. II. Psychological and phenomenological aspects. *Pain*, *45*, 239 – 248.
- Box, G., Jenkins, G. M., & Reinsel, G. (1994). *Time Series Analysis: Forecasting and Control*. Prentice Hall.
- Collins, F. S. (2010). The future of personalized medicine. *NIH Medline Plus*, *5*, 2–3.
- Collins, L. M., Murphy, S. A., & Bierman, K. L. (2004). A conceptual framework for adaptive preventive interventions. *Prevention Science*, *5*, 185–196.
- Deshpande, S. (2011). *A Control Engineering Approach for Designing an Optimized Treatment Plan for Fibromyalgia*. Master’s thesis Electrical Engineering, Arizona State University.
- Deshpande, S. (2014). *Optimal Input Signal Design for Data-Centric Identification and Control with Applications to Behavioral Health and Medicine*. Ph.D. thesis Electrical Engineering, Arizona State University, USA.
- Deshpande, S., Nandola, N. N., Rivera, D. E., & Younger, J. (2011). A control engineering approach for designing an optimized treatment plan for fibromyalgia. In *Proc. of the 2011 American Control Conference* (pp. 4798 – 4803).
- Deshpande, S., Rivera, D. E., & Younger, J. (2012). Towards patient-friendly input signal design for optimized pain treatment interventions. In *Proceedings of 16th IFAC Symposium on System Identification* (pp. 1311–1316).
- Deshpande, S., Rivera, D. E., Younger, J. W., & Nandola, N. N. (2014). A control systems engineering approach for adaptive behavioral interventions: illustration with a fibromyalgia intervention. *Translational Behavioral Medicine*, *4*, 275–289.
- Dong, Y., Deshpande, S., Rivera, D. E., Downs, D. S., & Savage, J. S. (2014). Hybrid model predictive control for sequential decision policies in adaptive behavioral interventions. In *Proceedings of the 2014 American Control Conference* (pp. 4198–4203).
- Garnier, H., & Young, P. C. (2014). The advantages of directly identifying continuous-time transfer function models in practical applications. *International Journal of Control*, *87*, 1319–1338.
- Gevers, M., Miskovic, L., Bonvin, D., & Karimi, A. (2006). Identification of multi-input systems: variance analysis and input design issues. *Automatica*, *42*, 559 – 572.
- Lee, J., & Yu, Z. (1994). Tuning of model predictive controllers for robust performance. *Comp. and Chem. Engg.*, *18*, 15 – 37.
- Lee, Y., Nassikas, N., & Clauw, D. (2011). The role of the central nervous system in the generation and maintenance of chronic pain in rheumatoid arthritis, osteoarthritis and fibromyalgia. *Arthritis Research and Therapy*, *13*, 211.
- Ljung, L. (1999). *System identification: theory for the user*. Upper Saddle River, NJ: Prentice Hall PTR.
- Ljung, L. (2009). Experiments with identification of continuous time models. In *Proceedings of 15th IFAC Symposium on System Identification* (pp. 1175–1180).
- Ljung, L., & Singh, R. (2012). Version 8 of the system identification toolbox. In *Proceedings of 16th IFAC Symposium on System Identification* (pp. 1826–1831).
- Martin, C. A., Rivera, D. E., Riley, W. T., Hekler, E. B., Buman, M. P., Adams, M. A., & King, A. C. (2014). A dynamical systems model of social cognitive theory. In *Proc. of the 2014 American Control Conf.* (pp. 2407–2412).
- Mattiloi, T. M., Milne, B., & Cahill, C. (2010). Ultra-low dose naltrexone attenuates chronic morphine-induced gliosis in rats. *Molecular Pain*, *6*, 1–11.
- Molenaar, P., & Campbell, C. (2009). The new person-specific paradigm in psychology. *Current Directions in*

Psychological Science, 18, 112–117.

- Nandola, N. N., & Rivera, D. E. (2013). An improved formulation of hybrid model predictive control with application to production-inventory systems. *IEEE Transactions on Control Systems Technology*, 21, 121–135.
- Perrot, S. (2008). Fibromyalgia syndrome: a relevant recent construction of an ancient condition? *Current Opinion in Supportive and Palliative Care*, 2, 122–127.
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11, 733–764.
- Riley, W., Rivera, D. E., Atienza, A., Nilsen, W., Allison, S., & Mermelstein, R. (2011). Health behavior models in the age of mobile interventions: are our theories up to the task? *Translational Behavioral Medicine*, 1, 53–71.
- Rivera, D. E., & Jimison, H. B. (2013). Systems modeling of behavior change: Two illustrations from optimized interventions for improved health outcomes. *IEEE Pulse*, 4, 41–47.
- Rivera, D. E., Pew, M. D., & Collins, L. M. (2007). Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction. *Drug and Alcohol Dependence*, 88, S31–S40.
- Timms, K. P., Rivera, D. E., Piper, M. E., & Collins, L. M. (2014). A hybrid model predictive control strategy for optimizing a smoking cessation intervention. In *Proc. of the 2014 American Control Conference* (pp. 2389–2394).
- Wang, W., & Rivera, D. (2008). Model predictive control for tactical decision-making in semiconductor manufacturing supply chain management. *IEEE Transactions on Control Systems Technology*, 16, 841–855.
- Wang, Y., Dassau, E., & Doyle, F. (2010). Closed-loop control of artificial pancreatic β -cell in type 1 diabetes mellitus using model predictive iterative learning control. *IEEE Trans. on Biomedical Engineering*, 57, 211–219.
- Wellstead, P., Bullinger, E., Kalamatianos, D., Mason, O., & Verwoerd, M. (2008). The role of control and system theory in systems biology. *Annual Reviews in Control*, 32, 33–47.
- Wolfe, F., Clauw, D. J., Fitzcharles, M.-A., Goldenberg, D. L., Katz, R. S., Mease, P., Russell, A. S., Russell, I. J., Winfield, J. B., & Yunus, M. B. (2010). The American College of Rheumatology preliminary diagnostic criteria for fibromyalgia and measurement of symptom severity. *Arthritis Care & Research*, 62, 600–610.
- Wolfe, F., Smythe, H. A., Yunus, M. B., Bennett, R. M., Bombardier, C., Goldenberg, D. L., Tugwell, P., Campbell, S. M., Abeles, M., Clark, P., Fam, A. G., Farber, S. J., Fiechtner, J. J., Michael Franklin, C., Gatter, R. A., Hamaty, D., Lessard, J., Lichtbroun, A. S., Masi, A. T., McCain, G. A., John Reynolds, W., Romano, T. J., Jon Russell, I., & Sheon, R. P. (1990). The American College of Rheumatology 1990 criteria for the classification of fibromyalgia. *Arthritis & Rheumatism*, 33, 160–172.
- Younger, J., & Mackey, S. (2009). Fibromyalgia symptoms are reduced by low-dose naltrexone: A pilot study. *Pain Medicine*, 10, 663–672.
- Younger, J., Noor, N., McCue, R., & Mackey, S. (2013). Low-dose naltrexone for the treatment of fibromyalgia: Findings of a small, randomized, double-blind, placebo-controlled, counterbalanced, crossover trial assessing daily pain levels. *Arthritis & Rheumatism*, 65, 529–538.
- Zafra-Cabeza, A., Rivera, D., Collins, L., Ridao, M., & Camacho, E. (2011). A risk-based Model Predictive Control approach to adaptive interventions in behavioral health. *IEEE Trans. on Control Systems Technology*, 19, 891–901.
- Zurawski, R., & Teel, A. R. (2006). A model predictive control based scheduling method for HIV therapy. *Journal of Theoretical Biology*, 238, 368–382.