Models, Devices, Properties, and Verification of Artificial Pancreas Systems.

Taisa Kushner[1], B. Wayne Bequette[2],
Faye Cameron[2], Gregory Forlenza[3], David Maahs[4] and Sriram Sankaranarayanan[1].

Abstract In this chapter, we present the interplay between models of human physiology, closed loop medical devices, correctness specifications and verification algorithms in the context of the artificial pancreas. The artificial pancreas refers to a series of increasingly sophisticated closed loop medical devices that automate the delivery of insulin to people with type-1 diabetes. On one hand, they hold the promise of easing the everyday burden of managing type-1 diabetes. On the other, they expose the patient to potentially deadly consequences of incorrect insulin delivery that could lead to coma or even death in the short term, or damage to critical organs such as the eyes, kidneys and the heart in the longer term. Verifying the correctness of these devices involves a careful modeling of human physiology, the medical device, and the surrounding disturbances at the right level of abstraction. We first provide a brief overview of insulin glucose regulation and the spectrum of associated mathematical models. At one end are physiological models that try to capture the transport, metabolism, uptake and interactions of insulin and glucose. On the end are data driven models which include time series models and neural networks. The first part of the chapter examines some of these models in detail in order to provide a basis for verifying medical devices. Next, we present some of the devices which are commonly used in blood glucose control, followed by a specification of key correctness properties and performance measures. Finally, we examine the application of some of the state-of-the-art approaches to verification and falsification of these properties to the models and devices considered. We conclude with a brief presentation on future directions for next generation artificial pancreas, and the challenges involved in reasoning about them.

1. University of Colorado, Boulder, USA. {taisa.kushner,srirams}@colorado.edu
2. Rensselaer Polytechnic Institute, Troy, USA. {bequette,camerf}@rpi.edu
3. Barbara Davis Center for Diabetes, University of Colorado Anschutz Medical Campus, Denver, USA. gregory.forlenza@ucdenver.edu
4. Stanford University Medical Center, Stanford, USA. dmaahs@stanford.edu

1 Introduction

This chapter presents research challenges pertaining to the application of automated reasoning and formal methods to the artificial pancreas project. The artificial pancreas is a *technological* treatment for type-1 diabetes, consisting of a set of devices which automate the external delivery of insulin to individuals with type-1 diabetes.

Type-1 diabetes is a condition that affects millions of individuals around the world. It is characterized by the destruction of the islet cells of the pancreas. These cells secrete the hormone insulin which is responsible for the regulation of blood glucose (BG) levels in the human body. In order to take up glucose from the blood, many cells, such as those of the musculoskeletal system, require insulin. As a result, a depletion or absence of insulin results in these cells being unable to take up glucose from the blood, leading these cells to starve even when BG levels are high. The treatment for type-1 diabetes is to replace the insulin secreted by the pancreas through external infusion of artificial insulin, known as *insulin analogs*. However, this insulin must be delivered in careful doses to keep the BG levels in a narrow *euglycemic* range of [70,180] mg/dL. BG levels below 70 mg/dL results in a condition known as *hypoglycemia*, which carries risks of loss of consciousness, coma or even death. On the other hand, BG levels above 180mg/dL result in a condition called *hyperglycemia*, and pose long-term risks which include damage to the heart, kidneys, eyes, and peripheral nerves. In order for insulin to be delivered, it is often infused subcutaneously through an insulin pump - a device that is programmed to deliver a constant low rate of insulin, known as *basal insulin*, to supply insulin during endogenous glucose production by the liver. Additionally, the patient has to deliver a bolus of insulin in advance of the meal to cover the carbohydrates and proteins that they will consume. This is performed by estimating the carbohydrate consumption and using a linear factor to determine the amount of insulin.

Additionally, an individual with diabetes must periodically check their BG levels using a glucometer, or a continuous glucose monitor, and deliver a "correction" bolus if high BG levels persist. Due to changes in insulin sensitivity during exercise, fever, or high stress, the patient must take special care to turn off their pump during these periods. Therefore, the patient acts as a manual controller for managing their own BG levels. This has many drawbacks: (a) it places a heavy burden on the patient who must constantly pay attention to their BG levels as many as 9-10 times a day; (b) mistakes such as miscounting carbohydrates, failing to turn off the pump, forgetting to bolus for meals and incorrectly estimating the meal bolus relative to the meal time, risk serious health consequences of falling into hypo- or hyperglycemia; and finally (c) hypoglycemia that happens during sleep can lead to dangerous consequences such as seizures or death, requiring constant vigilance.

The artificial pancreas project seeks to partially or fully automate the delivery of insulin by combining a continuous glucose monitor (CGM) which periodically senses BG levels subcutaneously, an insulin pump that delivers insulin, and a closed-loop control algorithm that uses inputs from the CGM and the user to control BG levels to a target value (typically in the range $80-120$mg/dL). A schematic diagram is shown in Figure 2. The development of the artificial pancreas was originally envisioned by the JDRF (formerly the Juvenile Diabetes Research Foundation), to be developed in three generations as shown in Figure 1. The development of these devices has proceeded concurrently. In this conception, each stage is built upon the capabilities of the previous stages with more automation,

lessening user-required inputs at each stage. The road map was thoroughly revised in 2015 to reclassify insulin delivery beyond stage 3 into "insulin-only" and "multihormonal" control [89]. These generations are all under various stages of clinical evaluation: the Medtronic 670G, a "stage 3" hybrid closed loop system, obtained FDA approval in 2016 [74], and many other devices are nearing commercialization at time of writing.

In theory, BG levels can be controlled by insulin infusion similar to how the velocity of a vehicle is controlled during cruise control. Simply put, the cruise control maintains the velocity at a set target value by increasing the throttle whenever the velocity falls below the target set-point, or braking when too much throttle was applied. Unfortunately, numerous obstacles prevent us from realizing this principle for the case of insulin infusion. While the first component exists, insulin can be increased to reduce BG levels, the second component, the break, does not. Glucagon, the counter-regulatory hormone to insulin which increases BG levels, is not currently available in a shelf stable form. Thus, we imagine a car with an accelerator, but no breaks. However, withholding insulin has the effect of bringing BG levels back up due to the endogenous glucose production in the liver, making insulin-only control a feasible option.

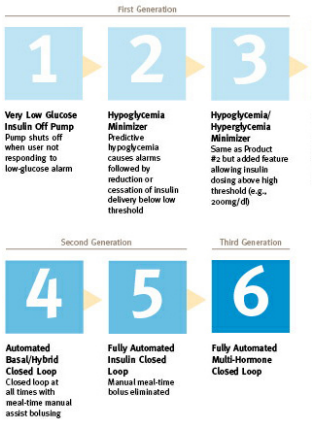

Stage:Refs.
1: [114]
2: [23]
3: [7, 68, 120]
4: [76–78]
5: [20–22, 33] [42, 87, 103]
6: [51, 52]

Fig. 1 AP Stages originally envisioned by JDRF, formerly the Juvenile Diabetes Research Foundation. This pathway was revised in 2015 [89].

In this insulin-only option, a serious limitation lies in the action profile of currently available insulin analogs. Insulin that is infused subcutaneously requires time to diffuse into the blood stream before it can begin to take action. Once in the blood, it achieves peak action in about 75 minutes, and remains in the system for nearly 3−5 hours before being metabolized. As a result, insulin infused at time t can remain in the system and continue to act for a much longer time than potentially needed. In the cruise-control analogy imagine a car with "sticky" gas pedal that take time to start acting but once it does, remains "stuck" for a while. Thus, the presence of delays and persistence of insulin with the absence of glucagon makes the problem quite hard. For instance, if a control responds to a surge in BG levels by pumping a proportional amount of insulin, this insulin will persist and continue to bring down BG levels well after the surge has vanished, exposing the patient to the risks of hypoglycemia.

Finally, user actions such as meals and physical activities have large and long lasting effects on the BG levels. Digestion of meals will result in a surge of glucose into the blood stream that must be counteracted by a dose of insulin delivered before the meal so that

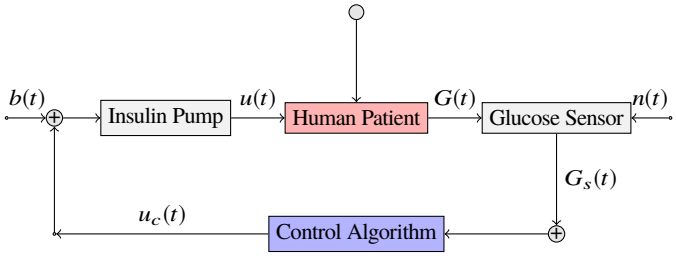

Fig. 2 Overview of the key components of an artificial pancreas control system. $b(t)$: external insulin, $u(t)$: insulin infused, $G(t)$: BG level, $n(t)$: measurement error, $G_s(t)$: sensed glucose level, $u_c(t)$: insulin infusion commanded.

the peak action of the delivered insulin coincides with the peak appearance of the meal carbohydrates in the blood stream which could be as late as 5 hours, depending on the type of food. Also, physical activity can lower BG levels and have long lasting effects well after the activity has abated. Unfortunately, unless specifically announced to the device, there are no currently available direct sensors which can be used to accurately warn the system of impending meals/exercise.

Thus, AP devices are, by the very nature of the problem, complex and safety critical. They need to be used by patients 24/7/365 without expert supervision, though they are capable of serious harm to the patient. As a result, their design and implementation requires careful consideration and thus form an ideal target for formal methods/automated reasoning approaches.

Structure of the Survey: In this survey, we will provide an overview of five important aspects of artificial pancreas systems. First, we will discuss human physiology, in particular the regulation of BG levels through insulin. Next, we discuss mathematical models of human physiology which form the basis for designing control algorithms. We discuss two important classes of control algorithms including PID and MPC algorithms with a focus on algorithms that are either in advanced stage clinical trials, or commercial products. Next, we discuss correctness and performance specifications. Finally, we discussion the application of verification techniques. The survey concludes by summarizing important challenges and the direction taken by next generation devices.

Comparison with Other Surveys: There have been numerous surveys and methodological reviews that have focused on various aspects of the problem of managing type-1 diabetes ranging from surveys that explain the nature of the problem including the capabilities of various technologies behind the artificial pancreas and the challenges underlying the development of control algorithms [35, 75]. More recent surveys including those by De Nicolau et al [119] and Lunze et al [101] have focused on the development of control algorithms. The lower level details of implementing model-predictive controllers for the artificial pancreas is described by Zavitanou et al [152]. In particular, they provide detailed descriptions of solver selection for MPC, the design of safety layers, a detailed hazard analysis, and also mention security risks. In contrast, this survey presents the clinical background as well as a description of the control algorithms to motivate the

challenges specific to the verification and personalization of these algorithms. Rather than focus on system level hazards, our focus is on the fundamental correctness of the core control algorithm in terms of whether it conforms to a detailed specification. We describe these issues in light of the recent developments, such as FDA approval of the first commercially available device in the US, the Medtronic 670G [74]. Additionally, we mention review articles that have argued in favor of specific approaches to AP control including Bequette's survey making the case for MPC [12] and Steil's review arguing the case for PID control [146]. An important review by Doyle et al surveys various control algorithms based on the clinical outcome measures and the need for improved head-to-head comparisons based on common measures of performance used in clinical evaluations [48].

2 The Glucose-Insulin Regulatory System

We first review the basic facts about the glucose-insulin regulatory system. Glucose is the sole source of energy to maintaining functioning of the brain and red blood cells, and is utilized to an extent by almost every other cell in the body. It is primarily obtained by digesting complex carbohydrates, which are then broken down into smaller molecules and enter the blood stream via the gut. However, food is consumed sporadically and blood glucose concentration must be maintained within a narrow range. Blood glucose concentrations which are too high can disrupt the osmotic balance in the human body and lead to long-term vascular disease. On the other hand, low blood glucose concentrations can lead to cellular starvation and neurological complications such as seizure. Thus, after each meal, the potentially large influx of glucose must be carefully regulated to be stored and transported to the cells as needed. The liver takes up excess glucose where it is stored for the short-term as glycogen. Glucose is also taken up by fat (adipose) cells where it is stored for the long-term as fat tissue. The stored glucose is released into the blood stream whenever needed through the breakdown of glycogen in the liver and the breakdown of fatty molecules in cases of extreme starvation.

In order for glucose to reach the cells and be used for energy, it must be transported throughout the body. This is accomplished via the blood stream, as dissolved blood glucose. Thus, in order to maintain a constant supply of energy to cells, the body constantly regulates blood glucose levels to be within the narrow range of about $70-110$mg/dL.

The two main hormones which regulate this process are produced by the pancreas: insulin and glucagon [5, 105, 112]. Insulin is a small peptide hormone produced by the β-cells of of the pancreas. It acts primarily on the insulin receptor to increase cellular intake of glucose into muscle and fat cells via translocation of the Glut-4 transporter. Glucagon is a peptide hormone produced by the alpha cells of the pancreas. It acts primarily on the Glucagon receptor to stimulate breakdown of glycogen and production of glucose in the liver. In this way, insulin and glucagon oppose each other to control blood glucose in the normal human body: in the presence of high blood glucose, insulin promotes moving glucose into the cells to lower blood glucose; in the presence of low blood glucose, glucagon promotes moving glucose out of the liver to raise blood glucose.

Rising blood sugar is sensed by pancreatic beta cells via passive diffusion of glucose into the cells through the Glut-2 transporter. When blood sugar levels rise, insulin production in the pancreatic β-cells is stimulated. Within about five minutes, the newly produced insulin crosses into interstitial space where it lowers blood glucose in three main ways:

1. Encourages most cells, especially muscle and fat cells to take up glucose (and other simple sugars) via the Glut-4 transporter
2. Inhibits breakdown of glycogen in the liver
3. Counters metabolic activity that would increase plasma glucose levels by inhibiting the conversion of animo acids or fats to glucose

This multi-delay feedback loop between glucose and insulin maintains healthy fasting blood glucose levels and controls for large disturbances such as meals and exercise, however it also results in a highly nonlinear, complex system which is difficult to remedy when a key component is off, such as in the case of diabetes [132].

3 Diabetes Mellitus

For some individuals, a metabolic disorder causes inhibitions in the body's production (hyposecretion), or use (hypoactivity) of insulin. Disorders of this type are broadly known as Diabetes Mellitus (or diabetes). The World Health Organization estimates that about 422 million individuals had a disorder of this type in 2014, and that number expected to rise to 642 million by 2040 making diabetes a defining issue in public health.

There are three main types of diabetes, Type 1, Type 2, and Gestational Diabetes. The most prevalent type of diabetes, *Type-2 Diabetes*, occurs in roughly 95% of cases. Individuals with Type 2 Diabetes are able to produce insulin, however the effects of insulin are deficient. *Type 1 Diabetes* results from the autoimmune destruction of the pancreatic β-cells, resulting in the individual being unable to produce insulin. The third type of diabetes, *Gestational Diabetes*, is a temporary condition which develops during pregnancy.

In all cases, blood sugar levels remain high after meals because glucose is unable to enter most tissue cells. The exceptions to this are liver, kidney, and brain tissues, as these cells are able to uptake glucose independent of insulin levels. When sustained, elevated blood sugar levels result in a condition known as *hyperglycemia*, which is a characteristic of diabetes. Typically, a diagnosis of diabetes is made when fasting blood glucose levels exceed 126mg/dL or when random glucose levels exceed 200 mg/dL [4].

Development of insulin-deficient hyperglycemia occurs in a paradoxical fashion whereby blood glucose is elevated, but intracellular glucose and energy molecules are deficient. Thus while there are excessive energy molecules available in the blood stream, there is energy deficiency within cells which are unable to uptake glucose without adequate insulin. In order to compensate, the body mobilizes fats which results in high levels of fatty acids and their metabolites, known as *ketones*, in the blood. Ketones are acidic, and their accumulation in the blood stream results in a drop in pH levels, triggering a condition known as *ketoacidosis*. If left untreated, ketoacidosis can disrupt oxygen transport, heart activity, trigger loss of electrolytes, depress the nervous system and even lead to death [27, 112, 142]. Even when treated, long-term complications of diabetes due

to persistent hyperglycemia include kidney disease or diabetic nephropathy, complications of retinopathy with potential blindness, nerve damage and cardiovascular disease [27, 142].

3.1 Treatment Strategies

Type-1 diabetes is treated by the injection of artificial insulin analogues either through a syringe or subcutaneously through an insulin pump. The use of external insulin is also common in insulin-dependent type-2 diabetes However, the latter occurs due to decreased sensitivity to insulin rather than a lack of insulin secretion. We will focus mostly on type-1 diabetes for the rest of this document. Treatments for diabetes depend both on the type and severity of the illness but in all cases include an interplay of diet, exercise and the individual having to 'be the pancreas": closely monitoring their blood glucose levels, and dosing insulin to maintain blood glucose in the healthy *euglycemic* range of 70-180mg/dL.

Unfortunately, as previously noted, the human glucose-insulin regulatory system is incredibly complex, highly nonlinear, and difficult to control. In addition to all of this, individuals with diabetes exhibit a large range of variations in physiological characteristics such as gender, weight, and insulin sensitivity, all of which affect their response to blood glucose levels and insulin doses. These characteristics vary not only between individuals, but within an individual over time as changes in weight, activity level, and growth and development during childhood, play major roles in altering insulin requirements throughout the lifetime.

As we noted, too little insulin is a problem. However, insulin levels which are too high are even more dangerous. Too much insulin will cause an individual to fall into *hypoglycemia*. Hypoglycemia causes glucose deprivation in the brain and central nervous system leading to brain damage [38]. Furthermore, complications from hypoglycemia occur on a faster time scale than those from hyperglycemia, and hence individuals must be very careful to not administer too much insulin. All of this makes the problem of self-regulation incredibly difficult and a large burden on the patient.

The AP system is currently the most promising treatment option for reducing patient burden and improving diabetes management. In order to manage the intricacies of BG control described above, a central component of the AP system is a model of human physiology used to predict future blood glucose levels. In the following section, we cover a history of such models, their contributions, as well as strengths and weakness of each.

4 Physiological Modeling

For several decades, scientists pursued a reductionist approach, moving from identifying events on the organ level, down to enzymes and genes. This enabled characterization of signaling pathways for how insulin binds to cells, enhances protein synthesis, and performs a variety of other functions [14, 92, 115]. However, in order to better understand the insulin-glucose response and identify how diabetes develops, a systems-level approach

was needed. This is achieved through use of mathematical models. These models have enabled novel insight into glucose-insulin dynamics, and provided an "artificial patient" test platform for evaluating diabetic treatments, in particular the AP system [126].

Since the early 1960s, numerous mathematical models have been proposed. In this section we survey a range of models and approaches, focusing on those which have resulted in key developments in the field. Mathematical modeling of the glucose-insulin system began with Victor Bolie in 1961. In order to retain the minimum possible mathematical complexity, Bolie employed a single compartment system using two first order differential equations to model glucose and insulin concentrations [18]. While his model was able to capture steady plasma glucose levels, it was limited to fluctuations in the range of 30-150mg/dL, and was unable to capture the oscillatory dynamics. On the other hand, Srinivasan et al. proposed a very complicated model. This model defined individual model states for lipid, carbohydrate, and protein loops, and included highly nonlinear interconnections [144]

The next pivotal breakthrough came in 1979 when Richard Bergman proposed what is now called the "minimal model" for insulin-glucose physiology. Through use of their ordinary differential equation (ODE) based model, Bergman and his coworkers were able to gain insight into blood glucose regulation while avoiding the hurdle of how to describe the β-cell [16]. To do this, they developed the intravenous glucose-tolerance test (IVGTT), and took frequent measurements of plasma glucose and insulin after glucose injection in dogs. This was both less invasive than the glucose-clamp and pancreatic suppression test methods of the time, and allowed them to reveal the complex dynamics between plasma glucose and insulin [13, 15, 16, 28].

The minimal model, presented in below, consists of two equations.

$$\frac{dG}{dt} = -(S_g + X(t)) \times G, \; \frac{dX}{dt} = p_2 \times I(t) - p_3 \times X(t).$$

Here G is plasma glucose, and X is insulin in the interstitial compartment. The first equation states that the effect of glucose (S_G), and the effect of interstitial insulin, ($X(t)$), act synergistically with glucose to restore plasma glucose levels after injection of glucose from meals. The second equation states that the flux of insulin from the plasma to the interstitial compartment is a product of the fractional rate of insulin appearance in the interstitial fluid, (p_2), with the plasma insulin, ($I(t)$), minus the fractional clearance of insulin from the interstitial compartment, (p_3), times the insulin in the interstitial compartment, ($X(t)$).

Since its publication, the minimal model has been validated by several other laboratories, and resulted in multiple major discoveries [49, 86]. Most importantly, they found that the effect of insulin is slow - insulin has to move from blood plasma to another compartment, later identified to be the interstitial fluid, in order to effect glucose.

While the minimal model is able to reasonably capture the nonlinear insulin-glucose response, it does not model aspects such as endogenous glucose production by the liver, insulin dependent vs. insulin independent uptake of glucose by various tissues in the body, and the effect of renal clearance of glucose which happens during hyperglycemia. Additionally, the model assumes that insulin and glucose inputs are given directly into the bloodstream, leaving it unable to model meals.

4.0.1 The Hovorka Model

Following the minimal model, more detailed physiological models were proposed, notably those of Hovorka et al. [76, 151] and Dalla Man et al [40, 108]. Each advancement in modeling was accompanied by novel biological studies, stressing the importance of how novel insight is gained through cross-disciplinary work.

The model proposed by Dr. Roman Hovorka and his colleagues builds upon the minimal model by introducing additional compartments. This Hovorka model separates glucose into two main compartments: an accessible glucose compartment, which represents blood plasma as well as tissues which are able to quickly equilibrate with plasma, as well as an inaccessible compartment which represents interstitial and intracellular space [76]. A single insulin compartment is used. These main compartments of the model are further broken down into effect compartments.

The model also includes insulin-independent glucose utilization, which is represented as a constant outflow from the accessible glucose compartment. Overall, the model contains nine equations and twelve parameters, which are theoretically identifiable, though require invasive studies with glucose tracers [76].

By separating insulin-dependent uptake of glucose from insulin-dependent uptake, Hovorka et al. were able to identify new nonlinearities in the glucose-insulin regulatory system. They found that, under fasting conditions, insulin-independent uptake dominates, and insulin-dependent uptake represents only about 13% of total glucose turnover. However, after meals, in a phase referred to as *postprandial* conditions, insulin-dependent uptake dominates. Hence an increase in basal insulin by 50% will have drastically different effects depending on when it occurs. While this 50% increase in basal insulin results in a 50% increase in insulin-dependent glucose uptake, the whole body glucose uptake changes by 50% only in the postprandial condition. During fasting conditions, the effect on whole body glucose uptake is only 7% [76, 77].

4.0.2 The Dalla Man Model and UVA/Padova Simulator

In 2006, Dalla Man et al. proposed a novel simulation model of the glucose-insulin system with mixed-meal simulation which was developed using data from both non-diabetic individuals, as well as those with type 2 diabetes [40, 108]. Prior to this, a few models describing oral rather than intravenous glucose absorption were available, however they were limited as validation of these models was performed on plasma glucose concentrations alone [95].

Aided by data from a triple-tracer glucose absorption study [9], Dalla Man et al. developed a simulation model which describes the physiological events that occur after a meal, and provides the basis for the FDA approved in-silico testing system, the "UVA/Padova simulator". This simulator has replaced animal trials and greatly reduced costs of preclinical testing [88, 126].

Similar to the minimal and Hovorka models, the Dalla Man model is a nonlinear ordinary differential equation model. It consists of 10 state variables with nearly 40 patient parameters, and is divided into subsystems. The two main subsystems of the Dalla Man

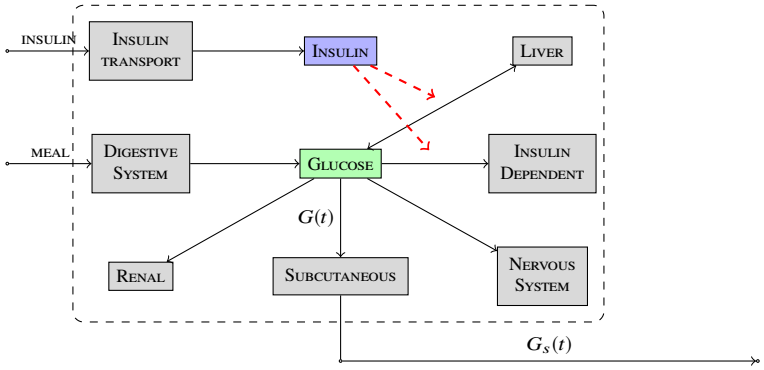

Fig. 3 Structure of the interaction between subsystems in the Dalla-Man et al model.

model are the glucose and insulin systems which are linked to one another via the control of insulin on glucose utilization as well as endogenous glucose production, EGP.

Figure 3 shows the overall structure of the model. The glucose subsystem of the Dalla Man model consists of two compartments, one for plasma glucose and rapidly equilibriating tissues, G_p, and one for the slowly equilibriating tissues, G_t both with units (mg/kg). System equations are presented below,

$$\dot{G_p}(t) = EGP(t) + Ra(t) - U_{ii}(t) - E(t) - k_1 \cdot G_p(t) + k_2 \cdot G_t(t)$$
$$\dot{G_t}(t) = -U_{id}(t) + k_1 \cdot G_p(t) - k_2 \text{cot} G_t(t)$$
$$G(t) = \frac{G_p}{V_G}$$

Here Ra is the rate of appearance of glucose in blood plasma (mg/kg/min), EGP is endogenous glucose production (mg/kg/min), U is utilization (mg/kg/min), G is plasma glucose concentration (mg/dL), V_G is glucose distribution volume (dL/kg), U_g and U_i are insulin independent and dependent glucose utilization, respectively, (mg/kg/min), k_1,k_2 are rate parameters (min^{-1}), and E denotes renal excretion (mg/kg/min) which occurs if plasma glucose exceeds a threshold:

The Dalla Man model also uses two compartments to describe insulin kinetics. The first equation, $I_l(t)$, denotes insulin masses in the liver, and the second equation, $I_p(t)$, denotes insulin concentrations in the blood plasma. Both have units (pmol/kg).

$$\dot{I_l}(t) = -(m_1 + m_3(t)) \cdot I_l(t) + m_2 I_p(t) + S(t)$$
$$\dot{I_p}(t) = -(m_2 + m_4) \cdot I_p(t) + m_1 \cdot I_l(t)$$
$$I(t) = \frac{I_p}{V_I}$$

Here I is plasma insulin concentration (pmol/L), S is insulin secretion (pmol/kg/min), V_I is insulin distribution volume (L/kg), and m_1,m_2,m_4 are rate parameters with units (min^{-1}).

Peripheral degradation, m_4 is assumed to be linear, and m_3 is related to hepatic extraction.

For the subsystems above, EGP,Ra, and U_g,U_i are the most important model processes. The Dalla Man model assumes that suppression of EGP is linearly dependent on the concentrations of plasma glucose (G_p), portal insulin (I_{po} with units pmol), and a delayed insulin signal (I_d with units pmol/L):

$$EGP(t) = k_{p1} - k_{p2} \cdot G_p(t) - k_{p3} \cdot I_d(t) - k_{p4} \cdot I_{po}(t) \quad (1)$$

here k_{p2},k_{p3},k_{p4} are parameters which govern liver glucose effectiveness (min^{-1}), amplitude of delayed insulin action (mg/kg/min per pmol/L), and amplitude of portal insulin action (mg/kg/min/pmol), respectively. The rate of glucose absorption, Ra is assumed to depend on two stomach compartments (solid and liquid phases), a gut compartment, and a rate constant for gastric emptying which is defined to be a nonlinear function of the amount of glucose in the stomach, identified in an earlier study [107].

Insulin independent glucose utilization, U_g is constant and represents glucose uptake by the brain and erythrocytes. Insulin dependent glucose utilization, U_i occurs in a remote compartment and depends nonlinearly on glucose in peripheral tissues.

$$U_g(t) = [2 \cdot \frac{PCR_b}{V_G} - S_G - (\frac{PCR_b}{V_G} - S_G) \cdot \frac{G_p(t)}{G_{pb}}] \cdot G_p(t)$$

$$U_i(t) = K_i(t) \cdot G_t(t) = X(t) \cdot G_t(t)$$

Here PCR_b is the basal plasma clearance rate in (dL/kg/min), S_G is fractional glucose effectiveness (min^{-1}), X is insulin action (min^{-1}), and the suffixes b denote basal state. These unit processes, EGP,Ra,U_g,U_i, were identified with a forcing function strategy, mathematical details for which can be found in [40, 108].

The full Dalla Man model has been incorporated into a simulator called the UVA/Padova Simulator, which provides a platform of virtual, *in silico* patients. This model, and its subsequent implementation into an easy-to-use simulator, has enabled the testing of safety and limitation of closed-loop AP systems, guided clinical studies, and has been a cost-effective method for pre-clinical testing. [32, 41, 69, 111] In 2013, the UVA/Padova simulator was updated to include the notion that insulin-dependent utilization of glucose increases nonlinearly when glucose falls below a selected threshold, and the generation of "virtual patients" was updated [109].

In addition to the aforementioned models, a variety of others have been proposed which utilize ordinary-, partial-, and even delay- and integro- differential equations [85, 97, 98, 121, 125].

Patient-Specific Models: While these models have been incredibly useful, a key critique is that these models are incredibly difficult to adapt to individual patients using readily available data. While a set of 300 virtual patients with varied demographics exists for use in the UVA-Padova Simulator, extending this to "real" patients is an ill-posed problem. The Dalla-Man model involves upwards of 40 patient parameters, 27 of which are free parameters. Some of these parameters can vary in time, some require intrusive radiological tracer studies to identify, and some have no biological relevance, making parameter estimation an ill-posed problem [34]. Recent work has begun to use Bayesian parameter inference methods for parameter estimation, however, fitting parameters remains an open area of research [128].

4.1 Data-Driven Approaches

The recent development and wise-spread use of continuous glucose sensors and insulin pumps has resulted in an upsurge of available data, and, consequently, has spurred the development of data-driven models. As with ODE-based models, a variety of data driven models and approaches have been proposed [64, 67, 70, 91, 123]. Broadly, these data-driven approaches seek to improve diabetic treatment protocols by characterizing glucose-insulin dynamics on the individual person level using readily available data that includes periodic BG levels and insulin infused. Whereas ODE-based approaches seek to develop models based on experimental knowledge of small molecule, enzyme, and tissue interactions, data-driven approaches seek to learn dynamics from data sets, most often time series data. However, it should be noted that ODE-based models can also be made "data-driven" by identifying parameters for the model that best fits the observed data. In fact, the UVA/Padova simulator with "virtual patients" falls under such an umbrella [34, 70, 77, 138].

Multiple papers have employed *auto-regressive moving average*, ARMA, based approaches [67, 91]. In general, given a time series of data points, an ARMA model provides a tool for understanding dynamics and predicting future values.

Given a discrete-time process with past states $x(t),x(t-1),...$, and inputs $u(t),u(t-1),...$, an ARMA model has the form:

$$x(t+1)=\sum_{i=0}^{p-1}a_i x(t-i)+\sum_{j=0}^{q-1}b_i u(t-j)+e(t),$$

wherein $a_0,...,a_{p-1}$ represent the coefficients of the "auto-regressive" part of the model, $b_0,...,b_{q-1}$ represent the "moving average" part of the model and $e(t)\simeq 0$ is the prediction error. Such models enable accurate, short term prediction, however they assume linearity and hence one needs to be cautious when applying them to prediction of nonlinear dynamics such as those of the glucose-insulin regulatory system.

A recent approach by Kushner et al [91] sought to overcome this limitation by intersecting multiple linear models to predict bounds on future BG levels (Fig. 4). These models were successfully used to optimize the tuning of gains, and verification of AP controllers for individual patients.

Statistical criteria, such as the partial autocorrelation and Akaike Information Criterion, were used to determine important time-lags in insulin action and predictors for future glucose levels. Based on these identified features, a parsimonious ARMA-like model was developed with order $p=2$ [3, 71, 91, 99] In order to include confidence bounds for BG prediction, an uncertainty interval, $[L,U]$, was added which resulted in the following nondeterministic model form which can predict BG levels with 99% confidence bounds:

$$G(t+\Delta_G)\in a_0 G(t)+a_1 G(t-\Delta_G)+bI(t-\Delta_I)+[L,U]$$

These models state that future glucose at $G(t+\Delta_G)$ can be predicted using the right hand side, using L as the lower confidence bound, and U as the upper bound.

To account for nonlinearity, five models fit with varying combinations of Δ_G,Δ_I, are intersected in a manner shown in 4. Intersecting models does present a trade-off with prediction accuracy, however, by Boole's inequality this new composite model prediction is at least 95% accurate.

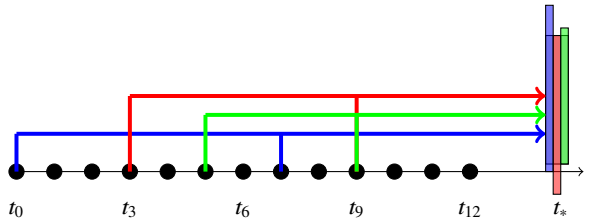

Fig. 4 Figure illustrating how multiple ARMA-based models are intersected to determine improved bounds for BG prediction. Each model uses a different prediction horizon, Δ_G, and pulls different combinations of inputs, Δ_I. These combinations are taken in a way such that all model predictions intersect at one future point in time, t_*. The final predicted range at this point is taken to be the intersection of ranges predicted by each model. Individual predicted ranges are depicted as vertical lines at t_* [91].

By using these models as "virtual patients" and coupling them with an *in silico* implementation of a FDA approved AP controller, Kushner et al. optimized the tuning of the control scheme, and increased time-in-range (blood glucose between 70-180mg/dL) for 83% of patients [61, 74, 91, 137, 145].

A shortcoming of the model by Kushner et al. is that it was not developed for use with disturbances such as meals and exercise. With respect to the development of AP control systems, a significant hurdle has been how to handle these two major disturbances: unannounced meals, and exercise. Consequently, a large number of models have been developed that try to characterize glucose-insulin dynamics in light of meals and exercise [41, 83, 123, 136].

Recent work by Paoletti et al [123], used data-driven methods to learn *uncertainty sets* from historic meal and exercise patterns in order to eliminate the need for meal announcements by the patient. An *uncertainty set* refers to the set of values that some parameter can take. To create a virtual patient, Paoletti et al used an extended version of the Hovorka ODE model for their "virtual patient" which they coupled with a model predictive control scheme [83, 136]. In order to identify meal and exercise schedules, Paoletti et al. learned *uncertainty sets* following a methodology developed by Bertsimas et al [17]. These uncertainty sets tried to capture patterns in when an individual would eat, or exercise, and were optimized against worst-case scenarios. Then, these uncertainty sets, along with the virtual patient model, were used to inform their AP controller of when an individual was likely to eat or exercise. The group showed their approach was able to keep virtual patients time-in-range around 93% [123].

Within the past year, data-driven approaches based on machine learning techniques such as neural networks, clustering, and support vector regression have gained traction [50, 57, 62, 63, 84, 127, 131]. Most approaches have predicted 30-60min horizons, with varying degrees of accuracy. A notable approach by Dutta et al. [50] utilized multiple neural networks to predict quantile bounds on future blood glucose levels. This approach is further discussed in 5.2.3. These approaches show promise, however they are not an end-all, and care should be taken to avoid increasing model complexity while yielding minimal improvement over linear predictions [131].

5 Control Algorithms

We briefly survey the vast space of control algorithms, focusing on two important class of algorithms for the artificial pancreas: PID-based control approaches and model predictive control (MPC) approaches. The design of control algorithms must tackle numerous challenges: (a) *time lags* in the sensor (10−20 minutes), the delayed action (15−20 mins) of subcutaneously delivered insulin and the persistence of insulin delivered in the body (3−4 hours) [28]; (b) large disturbances due to meals, physical activity, sensor noise, calibration errors, and infusion failures; and (c) one-sided control authority: insulin can reduce BG levels but the controller lacks the means to directly raise BG levels. Note that bihormonal control algorithms employ both insulin and glucagon [52]. However, stable form of glucagon is currently not available for use in a closed-loop AP although development is underway.

The AP controllers tested through clinical trials share many similarities in the control algorithms used: Proportional-Integral-Derivative (PID) control [146, 150], Model-Predictive Control [12, 21, 33, 68], and rule-based "fuzzy logic" control [48, 120]. In addition, *safety rules* which limit *insulin-on-board* - a term describing how much insulin is currently active in the body, handle pressure induced sensor attenuation (PISA) events [10], detect infusion set failures [79, 80] and request user intervention, are integral components of the control algorithms [134]. Cameron et al have proposed a multi-model predictive approach (MMPPC) that uses probabilistic models to predict the likelihood of future disturbances such as meals, thus taking a step towards removing the requirement for meal announcement [22]. An important advancement since carbohydrate counting is very burdensome on patients, and many patients do not comply with standard meal bolusing strategies advised by their clinical team. Other groups have since integrated meal detection strategies into their AP platforms [149]. Variations on this core MMPPC AP system have been tested in both inpatient [24] and hotel settings on people with T1D [26, 56].

5.1 PID Control

PID stands for "proportional, integral and derivative" control, a widely used scheme that controls a single target quantity (BG levels) through a control feedback which is computed as a simple summation of three terms based on the *control error*: $e(t)$: $G(t)-G^*$, wherein G^* is the desired target BG level. A PID controller delivers insulin according to the formula

$$u(t): \quad \underbrace{K_p e(t)}_{Proportional} + \underbrace{K_i \int_0^t e(s)ds}_{Integral} + \underbrace{K_d \frac{de(t)}{dt}}_{Derivative}.$$

Here, K_p, K_i, K_d are the proportional, integral and derivative gains, respectively. This simple scheme has numerous drawbacks when it comes to delivering insulin.

Saturation: First, the approach does not limit the maximum amount of insulin deliverable to the patient, which is required to ensure their safety. Next, it is easy to see that $u(t)$ can go negative, signifying some means to "take away" already delivered insulin. We ordinarily

$$
\begin{aligned}
I_e(t) &= I_e(t-5)+(G(t)-G^*) && \leftarrow \texttt{integral error}\\
D(t) &= \tfrac{G(t)-G(t-5)}{5} && \leftarrow \texttt{derivative}\\
I_p(t) &= K_0 i_d(t-5)+K_1 I_p(t-5)+K_2 I_p(t-10) && \leftarrow \texttt{insulin on board}\\
r(t) &= \begin{pmatrix} K_p(G(t)-G_0)+K_i I_e(t)+ \\ K_d D(t)-\gamma I_p(t) \end{pmatrix} && \leftarrow \texttt{raw control output}\\
u_c(t) &= \begin{cases} 0 & r(t)\le 0 \\ r(t) & 0\le r(t)\le i_{max} \\ i_{max} & r(t)\ge i_{max} \end{cases} && \leftarrow \texttt{saturated control output}
\end{aligned}
$$

Fig. 5 PID-based control with insulin feedback equations that govern the insulin delivered to the patient $u_c(t)$ based on the glucose input $G(t)$. The control law is based on parameters highlighted in blue. Brief explanations for each equation is also shown.

lack such means (pending shelf-stable glucagon). To counter this, the control algorithm *saturates* the insulin delivery to an upper limit $u_{\max}$ and a lower limit of 0.

Insulin Action: A major drawback of this approach is the delayed onset and persistence of insulin, mentioned previously. Delivering insulin according to the "current" BG levels alone does not account either the past infusions which can "stack" large amounts of insulin in the system, or for the fact that delivered insulin will persist well into the future. This problem can have fatal consequences for the patient.

Finally, it is well known that the presence of saturation in PID control leads to *integrator windup*, a condition where the built-up integral error, from when the control input was saturated, takes a prohibitively long time to "unwind". This leads the controller to deliver maximum insulin during this period and potentially endanger the patient.

PID Control Scheme for Artificial Pancreas: Steil et al proposed a hybrid closed loop controller based on the PID control scheme for delivering insulin, while at the same time overcoming the key limitations due to saturation and insulin action noted above [137, 145, 146, 150]. This control algorithm has been validated by numerous clinical trials and is the basis of a commercially available closed-loop system [74].

Figure 5 shows the equations that govern the calculation of the current insulin input, $u_c(t)$, from the input glucose signal, $G(t)$. First the PID controller is operated in discrete time with periodic updates to the value of $u(t)$ at 5 minute (or 1 minute) time intervals. This leads the integral error, $I_e(t)$, to be calculated as a summation rather than an integral of the overall error. Similarly, the derivative $D(t)$ is computed as a simple difference between current and past values divided by the time period. In order to account for past insulin, the controller tracks a term, $I_p(t)$, that is updated at each step according to the insulin delivered over the past 5 minutes, $i_d(t-5)$, and the past values of I_p, up to two time steps ago. The term $r(t)$ calculates the "raw insulin" by summing up the PID terms and adjusting for insulin already present in the system.

The controller is governed by parameters that include PID gains K_p: the proportional gain, K_i: the integral gain and K_d: the derivative gain. Other gains include that of the insulin feedback term, γ. The controller also uses a maximum insulin delivery rate i_{max}. Finding suitable values for these parameters is an important problem. Weinzimer et al

propose a basic rule of thumb for setting the gains using the daily insulin requirement of the patient [150]. We present an analysis of the rule using data driven models in Section 7.

5.2 MPC Control

Model predictive control (MPC) is based on the use of a prediction model to forecast the effects of proposed manipulated input changes on the desired output. At each time step k, an optimization problem is solved to obtain the current and future manipulated input changes for a control horizon of M moves, based on an objective function evaluated over a prediction horizon of P time steps. A major advantage to this approach is that constraints can be imposed directly, particularly on the manipulated inputs. MPC largely arose through applications in the chemical process industry to multivariable systems with many constraints [39]. Application to these large-scale systems was not limited by computational power because the sample times were often large (scale of many minutes). While the optimization problem solves for a sequence of control moves (step k to $k+M-1$), generally only the first control move (at time step k) is implemented. At the next time step (k+1), a new set of measurements is obtained, the model is updated, and a new optimization problem is solved. How the model is updated at each time step has a major impact on the closed-loop performance; Muske and Badgwell [116] discuss approaches that result in the best performance for unknown input disturbances.

This optimization problem has an analytical solution when the model is linear and no constraints are imposed. For a linear model with constraints, efficient quadratic programming (QP) algorithms are often used to find solutions [60]. Nonlinear model predictive control is a broad area with many challenges and many solution methodologies [135].

5.2.1 Artificial Pancreas Applications of MPC

While a number of PID [146] and fuzzy-logic (FL) [7, 113] based approaches have been proposed for closed-loop control algorithms, MPC-based methods have been most widely applied in clinical trials. For an overview of the algorithms, see Bequette [12]. While many articles have reported simulation-based results, the focus of this review is on algorithms that have been used in human clinical trials. Additionally, our effort is directed to the delivery of insulin and not to multi-hormone approaches that include glucagon.

Hovorka et al. presented a nonlinear model predictive control approach based on a seven-state compartmental model updated at different time intervals using Bayesian parameter estimation [77].

In this system, blood glucose is sampled and insulin infusion is adjusted using a 15-minute sample time. To calculate the next insulin infusion rate, $u(t+1)$, two components are considered: the difference between target trajectory, $y(t+i)$ and predicted trajectory, $\hat{y}(t+i|t)$, and the variation in the infusion rate, $u(t)$. These are balanced by changing k_{arg}, the 'aggressiveness' constant of the controller such that this difference is minimized over a 4 hour prediction horizon:

$$\min_{0\le u(t+1)\ldots u(t+N)\le 4}\{\sum_{i=1}^{N}[\hat{y}(t+i|t)-y(t+i)]^2+\frac{1}{k_{arg}}\sum_{i=1}^{N}[u(t+i)-u(t+i-1)]^2\}$$

Clinical studies were performed on 15 subjects, using prediction and control horizons of 4 hours (16 time steps). The desired glucose reference trajectory depended on whether the current glucose was high or low: if high, a relatively slow trajectory is used, and if low, a rapid first-order approach to the setpoint is specified. These trials were post-dinner, overnight studies with the controller activated several hours after the last meal, and thus did not tackle the challenges of meals or exercise.

There are many advantages of using MPC in automated insulin dosing, as discussed by Bequette [12]. A number of closed-loop algorithms are directly compared in simulation studies by Cameron et al [21]. A clinical trial directly comparing MPC and PID is reported by Pinsker et al [129], who show that MPC has better performance (lower mean glucose and longer glucose time-in-range) in a 27.5-hour trial involving 20 subjects.

Tuning and Adaptation of MPC Algorithms: MPC algorithms are usually tuned according to an individualized subject parameter, such as insulin sensitivity (drop in blood glucose for each unit of insulin delivered), or the carb/insulin ratio (the inverse of the amount of insulin needed for each gram of carbohydrate consumed). These factors can also be estimated based on an individuals total daily insulin dose [82]. Models can also be adapted in real-time, as detailed by Turksoy et al [147, 148].

Feedforward Control: Meals, Exercise and Other Events. The majority of control algorithms require feedforward action for "events" that have a major effect on blood glucose and insulin requirements to maintain desired glucose. For example, the meal is ideally "announced" with an estimate of the amount of carbohydrates to be consumed, and an insulin bolus based on a known insulin/carb ratio. However, to reduce patient burden, it is desirable to have algorithms that can detect a meal based on changing glucose levels and other information. Explicit meal detection for improved MPC performance is presented by Lee et al [93, 94]. A probabilistic method for meal detection is presented by Cameron et al., who use a 7-state discrete compartmental model [22].

Exercise can have a rapid affect on blood glucose, so it is desirable to announce exercise and reduce insulin delivery well in advance of exercise. If an individual consumes a snack in anticipation of exercise, it is important that any meal detection features be deactivated to avoid risky insulin delivery. Activity monitors to detect exercise has been used in the algorithms presented by Cameron et al [26] and Turksoy et al [149].

State estimation: Whereas MPC approaches require the state of the predictive model at each time step, we are often unable to directly measure all the state variables used in the predictive model, and the measurements on actual patients are noisy or unsynchronized. For AP systems in particular, these measurements may come from CGMs, insulin pumps, and activity monitors. In order to clean these measurements and extract the components needed for the initial model state from which the optimization problem is calculated, a *state estimator* is often used. Many varieties of estimators such as Luenberger observers, Kalman filters, and Moving-Horizon Estimators. [19, 37, 65, 66]

Kalman filters, which use a series of measurements taken over time to smooth noise and other inaccuracies, have been particularly popular for estimating rates of change in glucose, as well as meal inputs [36, 93, 94] Methods which exploit device-specific

knowledge, such as sensor calibrations have also been successfully used. In particular, Gondhalekar et al. [65] used a moving-horizon like estimator to ameliorate effects of delays between measurements and controller action, instances of CGM sensor drop out, and discontinuities due to CGM calibrations.

5.2.2 Glucose Control Objectives

The vast majority of MPC algorithms use an objective function based on the mean glucose prediction over a future time horizon. Most often the objective function minimizes the sum of the squares of predicted errors (setpoint minus predicted output) over the prediction horizon. However, rather than controlling to a specific setpoint, the zone MPC approach attempts to keep BG within a desired range, over the prediction horizon. Harvey et al. studied 12 subjects in a 24-hour trial, using a zone of 80-140 mg/dL [73]. Forlenza et al compared zone MPC with sensor augmented pump therapy in a crossover trial involving 19 subjects [55]. While the zone MPC approach presented above is for day and night use, its performance is comparable to the predictive hypo/hyperglycemia minimizer used in overnight studies by Spaic et al. [143] and Bequette et al. [11].

The multiple model probabilistic predictive control (MMPPC) approach, on the other hand, manipulates insulin to achieve a probability of a hypoglycemic risk (2.5% probability of going below 100 mg/dL, for example) over the prediction horizon [22]. The approach is anticipatory, based on a database of common eating patterns. For example, as the time becomes longer between meals the probability that a meal will be consumed in the near future increase; similarly, when an individual has been awake for a longer period of time, the probability that they will go to sleep increase, resulting in less aggressive insulin delivery to avoid the risk of nocturnal hypoglycemia.

MMPPC results are presented by Cameron et al. for 10 subjects using a laptop-based platform in a clinical research center study [24]. Using an ambulatory platform, Cameron et al. [26] present results for 10 subjects in a 30-hour in-clinic study at two sites, followed by 15 subjects in 54-hour hotel-based study at three sites. Forlenza et al. study six adults and four adolescents in a 72-hour hotel-based study at three sites [56]. The study compared announced vs unannounced meals with 4-hour post-meal glucose averaging 140.6 and 197.8 mg/dL, for announced and unannounced meals respectively.

5.2.3 Robust MPC

In practice, all models used in MPC-based approaches have an inherent uncertainty due to inter- and intra-patient variability with respect to insulin sensitivity and glucose utilization, sensor noise, and meal detection. Multiple groups have taken up this problem of developing *robust* MPC approaches [50, 93, 104, 124]. Lee and Bequette [93] developed a deterministic multisine approach to estimate insulin-to-glucose dynamics and estimate meal size when tested *in silico* on virtual patients using the Hovorka model described in Section 4.0.1.

A novel approach developed by Dutta et al. [50] combines quantile prediction with neural network models to learn robust MPC controllers. By utilizing three separate patient models in their MPC approach, Dutta et al. were able to bound blood glucose predictions with 95% certainty. This resulted in controllers which provide model guarantees and bound risks on hypo and hyperglycemia for real data and simulated patients, and can handle unannounced meals.

6 Specifications

In this section, we describe correctness and performance specifications for artificial pancreas controllers. These specifications are of great importance since they ultimately define the means by which the patient's safety, and the efficacy of the treatment, can be guaranteed. We will also motivate the need to carefully model "usage scenarios" under which correctness properties can be asserted.

The most important correctness properties include "appropriate" behavior of the controller under hypoglycemia (defined as BG $\leq$ 70mg/dL) and hyperglycemia (defined as BG $\geq$ 180mg/dL). It is tempting to specify that these conditions should *never* occur in the first place. However, this objective is quite unlikely to be achieved by any algorithm, pending significant scientific/engineering advances such as shelf-stable glucagon that provides the ability to counteract delivered insulin; insulin analogues that act faster and persist for a shorter duration; and highly accurate CGM sensors devoid of dropouts, sensor attenuation or other types of failures. Fundamentally, the magnitude of disturbances due to meals, exercise, adrenaline response, and sensor errors can rapidly cause the patient's BG to fall below hypoglycemic limits or above hyperglycemic limits before the controller can sense/act upon these changes. Inter-patient variability is yet another major impediment to guaranteeing correctness. As we have seen thus far, control algorithms exhibit numerous parameters that must be carefully "tuned" to individual patients for correctness and performance. Individual variations are characterized roughly in terms of insulin sensitivity, BG responses to meals and physical activities, or through a large number of model parameters that are hard to identify from available data.

Thus, specifications must be coupled with usage scenarios which constrain user behavior such as the meals (timing, amount, and duration), exercise (timing and intensity), and accurate announcements of impending meals and physical activities. Furthermore, the scenarios themselves may be dependent on specific combinations of patient characteristics and control algorithm parameters. At the time of writing, the mapping between specific patient characteristics and algorithm parameters remains an open challenge.

6.1 Interface Correctness

One of the most fundamental correctness specifications addresses the *interface* between the algorithm, the insulin pump, and the CGM. First, the device must receive the data

from the CGM for the patient being treated without any form of data corruption. For instance, "cross-talk" between CGMs of two different patients in close physical proximity is a potentially dangerous situation. Likewise, CGM values can be "spoofed" by malicious attackers to cause serious consequences, including death to the patient [133]. One approach to defending against some of these issues is through standard cryptographic encryption and authentication. Another solution is to implement filters that can "smooth" away rapid fluctuations in CGM readings - scenarios which are potentially indicative of corrupted data.

The interface to the insulin pump is another safety critical component [2, 96]. First, the commands issued by the algorithm to the pump must be carried out in an accurate, reliable, and verifiable manner. Next, user inputs such as manual insulin boluses must be reported to the algorithm and taken into account for future insulin delivery. For instance, if a user-provided bolus is not reported to the algorithm, the *insulin on board* estimate diverges from the actual value, risking serious harm to the patient.

It is notable that many practitioners in the field, including the authors of this paper, recognize the issue of interface correctness. Furthermore, adverse events involving incorrect interfaces are known anecdotally from past clinical trials. However, this issue has received relatively little discussion in the academic literature at the time of writing. Most interfaces to CGMs and insulin pumps remain proprietary, with very few accepted standards across various commercially available models.

6.2 Algorithm Correctness and Performance

Algorithm correctness has many aspects including: (a) the algorithm must not command actions that make adverse situations worse – for instance, fail to shut down the pump during hypoglycemia, or shut down the pump when the BG level is too high; and (b) the algorithm must guarantee that under carefully defined nominal (reasonable) usage conditions, the patient should not exhibit hypoglycemia or extreme hyperglycemia.

As mentioned earlier, correctness in terms of avoiding hypo/hyperglycemia requires us to carefully define "usage" scenarios in terms of meals and physical activities. Figure 6 depicts a set of possible usage scenarios defined by ranges over the meal timings, meal amounts, the timings of boluses relative to the meal, and the timing of when the user switches from open loop to closed loop control mode. Other aspects of the scenario may include specifying carbohydrates counting errors, wherein the user reports the wrong amount of carbohydrate to the algorithm. Some of the parameters used in a previous work are reported in Table 1.

Correctness properties are specified using common logics such as Metric/Signal Temporal logics [90, 106]. Temporal logics are a widely used formalism across many areas of hardware and software verification. A temporal logic formula specifies the desired behavior of a system over time using logical connectives such as And ($\wedge$), Or ($\vee$), Not ($\neg$) and Implications ($\Rightarrow$) in addition to temporal connectives such as Always ($\mathcal{G}$), Eventually ($\mathcal{F}$) and Until ($\mathcal{U}$) [8, 110]. This allows assertions about the behavior of a system to be translated into a computer readable formalism. In addition, these constructs can be further

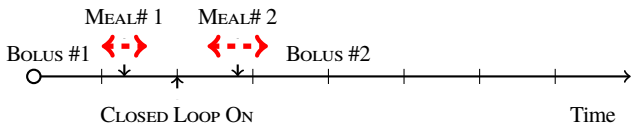

Fig. 6 Illustration of a usage scenario specification that includes two meals with limits on their timings and carbohydrates along with time range for switching to the closed-loop system.

enriched with time intervals that allow us to specify that certain events must or must not happen inside specific time intervals.

Let us consider an example of such a specification. We wish to specify that "whenever the blood glucose levels fall below 70mg/dL, the insulin pump must be shutoff, and remain so until the levels rise above 70 mg/dl and remain so for at least 10 minutes.

$$\underbrace{G \leq 70\text{mg/dL}}_{\text{BG levels below 70 mg/dL}} \Rightarrow \left(\underbrace{(u_I = 0)}_{\text{No insulin input}} \underbrace{\mathcal{U}}_{Until} \underbrace{(\mathcal{G}_{[0,10]} G > 70\text{mg/dL})}_{\text{BG is continuously above 70 mg/dL for 10 minutes}} \right)$$

Temporal logic specifications are a starting point for building safe and verified medical devices. Examples of complex temporal specifications for the artificial pancreas are described in previous work involving subsets of the coauthors that provided property specification for the PID control algorithm [25] and a Kalman-filter based pump shutoff controller [139].

Specifying Performance: A large variety of performance metrics have been studied for AP systems through various clinical trials [81]. Perhaps the most important performance metrics include the percentage of time that the BG levels remain inside the euglycemic range of [70,180]mg/dL. Studies that focus on longer term usage of AP systems have focused on the improvements to the HbA1C levels, a biochemical measure of the average BG levels over a period of time. However, the HbA1C levels are hard to simulate in mathematical models and are in fact a proxy for the average time inside the euglycemic range. An important class of performance metrics compare the behavior under closed loop against that under open loop. Such a comparison is hard to perform clinically since the patient's physiological state may vary between sessions. However, it is possible to do so in simulation since we can initialize the model to the same initial state. This class of performance metrics compares the insulin infused under a closed loop system against the daily insulin requirements by adding up the basal insulin and boluses.

As is the case with "correctness", performance requires specifying a distribution of usage scenarios. In-patient clinical trials control for this distribution by making the patients adhere to a strict meal and exercise protocol before and during the trial. The relatively shorter duration of these trials makes this feasible. However, controlling usage scenarios is hard for outpatient trials. For instance, it is well known that the use of a new device can change patient's meal and exercise habits either through positive or negative reinforcement. Currently, modeling patient's usage of devices and how they may change over time due

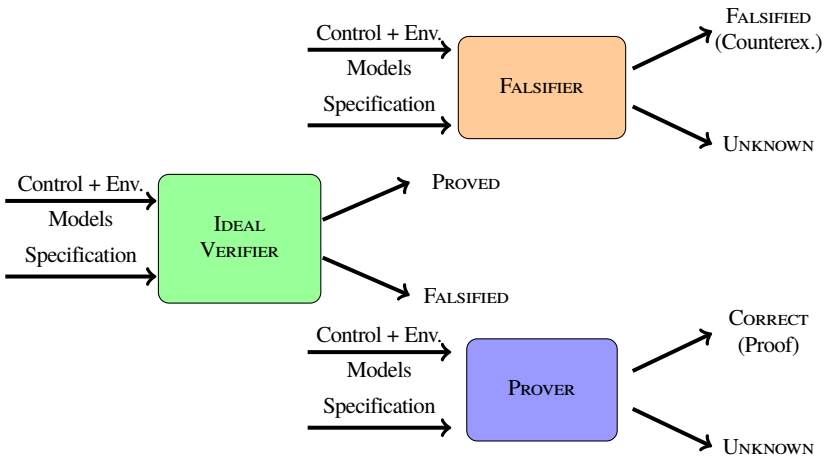

Fig. 7 An ideal verifier that attempts to prove as well as disprove at the same time is infeasible due to the undecidability of the halting problem. As a result, many approaches focus on "best effort" search for counterexamples (falsification) or finding proofs of correctness (verification).

to the introduction of a new AP closed loop system is an important unmet challenge for clinical researchers as well as engineers.

7 Verification and Synthesis

We provide a brief overview of formal methods that encompass mathematically rigorous approaches to the design and verification of safety-critical systems. Formal methods (along with computationally less expensive heuristic approaches known as semi-formal methods) consist of logical and machine readable specifications that enable the designer to write down correctness specifications for the system, modeling approaches that can capture the behavior of the controller and the operating environment; as well as verification approaches that establish proofs of correctness or find counterexample traces.

We have already looked at correctness specifications using formalisms such as temporal logic in the previous section. The verification problem takes as input (a) models of the system and its operating environment, and (b) the desired correctness specifications, and finds out whether or not the closed loop consisting of the system and environment satisfy the desired specifications. This is often achieved in practice through informal processes such as code review, unit /system tests, focusing on various forms of code coverage and finally, monitoring the system through deployment. In contrast, formal verification seeks to automate the verification process by using machine readable specifications of correctness, models of the control algorithm and its operating environment. Figure 7 shows a schematic of a formal verification procedure that takes in a control design, an environment specification, and correctness properties. It computes over these inputs and

declares that the system is correct w.r.t specifications (optionally providing a formal proof that can be machine checked), or a disproof in the form of environment inputs that cause the closed loop to violate a specification. Although a formal verifier is quite desirable to construct, there are fundamental limitations due to the undecidability of the halting problem [141] that prevent us from building a perfect verifier. Nevertheless, most verification approaches can be split into two main categories: falsification approaches that seek to discover counter-examples or proof approaches that seek to discover reasons why the system is correct.

Falsification approaches are based on exhaustive enumeration of all possible system behaviors (a process called *model checking* [8, 31]). Recently, the use of constraint solvers in model checking have enhanced the power of these approaches to search through an astronomically large number of possibilities, while converging onto violations rapidly [43]. However, existing constraint solvers are limited to restricted classes of constraints such as linear arithmetic with uninterpreted function symbols. AP systems are often nonlinear, and exhibit a mix of continuous and discrete behaviors that make them particularly challenging. Solvers for handling such systems are emerging at the time of writing and a subject of much research in the formal methods community [59]. Another class of falsification approaches rely on the concept of robustness of temporal formulas with respect to traces. A more detailed description is provided in the latter half of this section.

In contrast to falsification search, proof approaches presume that the system is correct and search for a mathematical proof of correctness. The proof itself can be written out as a sequence of logical inferences and represented as such in a computer. This approach is the basis of manual/semi-automated proof engines such as Keymaera [130]. However, these systems require expert user interaction to guide the process of proving. In particular, expertise is currently needed both in terms of understanding the design under verification and the symbolic logic/ proof techniques used by the theorem prover. This gap remains the focus on ongoing research in this area. Another approach called flowpipe construction relies on simulating the system over sets of possible states in order to predict the behaviors resulting from all these states in one shot [30, 58]. Currently flowpipe approaches have demonstrated the capability to handle large linear ODEs interacting with software-based control systems. However, their ability to handle nonlinear systems is an area of active research.

7.1 Verification of AP Algorithms

We will now survey existing approaches to verification, summarizing results from two representative case studies involving the PID-based system described in Section 5.1. AP controllers are safety critical and require a high level of assurance. Currently, simulations on virtual patient models and multiple stages of clinical trials on increasing numbers of patients under in-patient, hotel/camp and outpatient settings are used to provide assurance along with traditional software engineering practices such as design and code reviews. However, AP devices are used by patients who are not necessarily experts in human physiology or control systems. Furthermore, these devices will be expected to

T_1	[0, 60] mins.	Time for meal # 1
T_2	[180, 300] mins.	Time for meal # 2
X_1	[50, 150] gms.	# Carbs in meal # 1
X_2	[0, 40] gms.	# Carbs in meal # 2
IC_1, IC_2	[0,0.01] U/gm	Insulin to Carbs Ratio
δ_1, δ_2	[-15,15] min	Timing of insulin bolus relative to meal
$d(t), t \in \{100,105,\ldots,720\}$	[-20, 20]mg/dL	Sensor error

Table 1 Parameters defining user scenario for the simulation based falsification of the PID algorithm. The S-Taliro tool explores parameter values in the defined ranges to falsify given temporal properties over the closed loop system.

be operational 24/7/365, dealing with largely unanticipated disturbances due to user activities, physiological variations and component failures. Verification approaches that can exhaustively analyze the behavior of the algorithm under a numerous usage conditions are can help us evaluate safety and performance prior to large scale deployment.

A growing volume of approaches to verifying specific AP algorithms have been proposed thus far. Chen et al studied a PID-based closed loop system meant for intraoperative use in patients [29], using the dReal SMT solver [59] to prove safety for a range of parameters and controller gains. Other approaches to verifying AP algorithms have relied on *falsification*, using temporal logic robustness [47, 54], and incorporation of tools such as S-Taliro [1, 117] and Breach [46].

A subset of the coauthors have studied the use of falsification techniques for verifying closed loop control systems for the AP [25]. Their work investigated a PID controller described in section 5.1. Another study by the same team [139] was performed to test a predictive pump shutoff controller designed by Cameron et al. [23] that has undergone outpatient clinical trials [102, 143]. Recently, the PID controller in section 5.1 was studied over data-driven models using data from specific patients to synthesize gain parameters [91]. Shmarov et al present an approach to tuning PID controllers using stochastic hybrid system models to capture the patient and disturbances [140]. However, the approach focuses on tuning for a well defined ODE model of the patient physiology, whose advantages and drawbacks have been previously mentioned.

We will describe two different approaches to verification of the PID controller: (a) using simulation-based falsification applied to a detailed simulation model of the PID control algorithm [25]; and (b) using a bounded-model checking approach applied to a data-driven patient model described in section 4.1 and the PID control algorithm [91].

7.2 Simulation-Based Falsification of PID Algorithm

Sankaranarayanan et al performed a simulation-based verification of the PID control algorithm using a detailed simulation model that combines the Dalla Man model for the patient's insulin-glucose regulation, the control algorithm implemented from published descriptions [137, 145, 150] and constraints specifying the patients meals and boluses,

Property	Result
BG levels always ≥ 70mg/dL	NOT FALSIFIED
BG levels always ≤ 300mg/dL	FALSIFIED
Insulin shutoff whenever BG ≤ 90mg/dL	FALSIFIED
No hyperglycemia 10 hrs after start of simulation	NOT FALSIFIED
No hyperglycemia lasting more than 3 hrs	FALSIFIED
No hypoglycemia lasting more than 150 mins	NOT FALSIFIED

Table 2 Falsification results for various properties of the PID algorithm, as originally reported by Cameron et al [25].

as described in section 6 and detailed in Table 1. These components were implemented using Matlab(tm), and the tool S-Taliro developed by Fainekos et al was used to drive the falsification [1, 6]. Given a simulation model and a property of interest expressed in metric temporal logic [90], S-Taliro systematically searches for behaviors that minimize the robustness of the model's traces with respect to the specification of finding a trace that falsifies the specification. To generate traces, the tool adjusts the various model inputs including (a) the parameters affecting the usage scenario in terms of meal amounts, timings, and insulin boluses (see Table 1); and (b) the discrepancy/error between the underlying BG values in the simulation model and the values reported to the algorithm. Normally, this error has a stochastic nature [53], but is treated as an "adversarial" input to search for violations. S-Taliro employs a variety of stochastic search algorithms including random search, simulated annealing and genetic algorithms. The result of the approach either finds a falsifying trace that violates the property, or fails to produce a trace after an upper limit on the number of simulations has been exhausted. The properties examined and the results from S-Taliro are summarized in Table 2.

7.3 Verification Using Data-Driven Models

In this section, we discuss an approach used for verification and improved tuning of the PID control reported in Kusher et al [91]. The main specification is that the user's BG levels remain in a "safe" (euglycemic) range of $70-180$mg/dL. The approach combines multiple linear data driven models, as described in Section 4.1. The models were derived from overnight BG data for nearly 50 patients with 40 nights of data for each patient. For each patient, multiple models are derived with each model predicting a future BG value $G(t+T)$ as a linear function of past BG values, insulin on board values, and a non-deterministic uncertainty.

$$G(t+T) \in \sum_{j=1}^{p} a_j G(t-T_j) + \sum_{j=1}^{q} b_j I(t-T_j) + [\ell, u].$$

Multiple models were selected for various values of future lookahead T and past delays T_j for $j = 1,\ldots,p$. The models are treated as "data-driven" constraints that restrict the range of future BG values. The insulin inputs are provided to the patient using a closed loop algorithm described in Section 5.1. In order to determine whether an individual

Real	$G(-\Delta_{G,max}),\ldots,G(-5)$	History Gluc.
	$\mathrm{I}(-\Delta_{I,max}),\ldots,\mathrm{I}(-5)$	History IOB.
	$G(0),G(5),\ldots,G(T)$	Glucose
	$\mathrm{I}(0),\mathrm{I}(5),\ldots,\mathrm{I}(T)$	IOB
	$u_c(0),u_c(5),\ldots,u_c(T)$	Insulin
	$I_e(0),I_e(5),\ldots,I_e(T)$	Integral Err.
	$I_p(0),I_p(5),\ldots,I_p(T)$	Ins. feedback
	$r(0),r(5),\ldots,r(T)$	Raw ins.
Binary	$w(0),w(5),\ldots,w(T)$	0-1 variables
maximize $G(T)$ minimize for lower bnd		
$G_{min} \le G(t) \le G_{max}$		$t<0$
$\mathrm{I}_{min} \le \mathrm{I}(t) \le \mathrm{I}_{\max}$		Historical range
$G(t+\Delta_{G,i}) \le a_{0,i}G(t)+a_{1,i}G(t-\Delta_{G,i}+b_i\mathrm{I}(t-\Delta_{I,i})+u_i$		
$G(t+\Delta_{G,i}) \ge a_{0,i}G(t)+a_{1,i}G(t-\Delta_{G,i}+b_i\mathrm{I}(t-\Delta_{I,i})+l_i$		
for all $t=0,5,\ldots,T, i=1,\ldots,k$		Model M_i constr.
$\mathrm{I}(t)=F(u_d(t),u_d(t-5),\mathrm{I}(t-5),\mathrm{I}(t-10))$		Insulin On Board Equation
$I_e(t)=I_e(t-5)+(G(t)-G_0)$		Integral error
$I_p(t)=K_0 i_d(t-5)+K_1 I_p(t-5)+K_2 I_p(t-10)$		IOB feedback
$r(t)=\begin{pmatrix} K_p(G(t)-G_0)+K_i I_e(t)+ \\ K_d \frac{(G(t)-G(t-5))}{5}-\gamma I_p(t) \end{pmatrix}$		Raw ctrl. output
$u_c(t) \ge r(t)$		
$u_c(t) \le r(t)+i_{max}w(t)$		
$u_c(t) \le i_{max}(1-w(t))$		
$u_c(t) \ge 0$		Saturation

Fig. 8 Optimization problem for finding the maximum possible $G(T)$, wherein the constraints describe an execution of the closed loop according to the data-driven models $M_1,\ldots,M_k$ and PID control equations in Fig. 5 modified by encoding the saturation term using 0-1 variables.

will be safe under the action of the device, the maximum and minimum blood glucose values attainable under all possible model and controller behaviors are calculated using a mixed-integer optimization problem as shown in Fig. 8. This problem encodes the behavior of the model in terms of constraints, and the relation between the BG values output by the model and the insulin inputs of the controller. The integer variables arise from the need to model saturation of the insulin delivered by the PID controller.

In order to encode this loop, each patient model is to identify upper and lower constraints on future blood glucose levels. Next, the control scheme from [145, 150] is incorporated as constraints on insulin input.

This analysis requires modeling the patient's behavior prior to the start by assuming a "reasonable" history of blood glucose and insulin inputs. The following history assumptions are used,

1. For all $t<0$, $G(t) \in [G_{min},G_{max}]$ wherein $G_{min}=70$mg/dL and $G_{max}=180$mg/dL
2. For all $t<0$, $\mathrm{I}(t) \in [\mathrm{I}_{min},\mathrm{I}_{max}]$ wherein $\mathrm{I}_{min}=0.5$ U and $\mathrm{I}_{max}=4.0$ U

As such, these assumptions capture an infinite set of possible histories before $t=0$, when the controller is turned on.

The verification is run over a time horizon T with the control law run every $\Delta=5$ minutes. At each time step, the worst-possible attainable high and low blood glucose

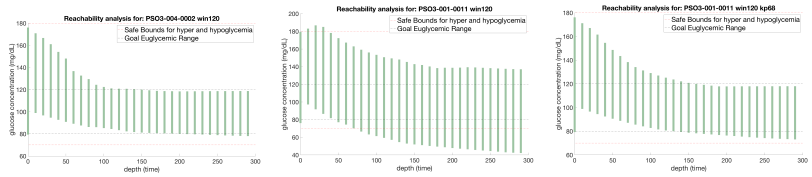

(a) **(b)** **(c)**

Fig. 9 Reachability analysis plots showing ranges of glucose values across time for selected individuals. Green bars denote all possible blood glucose values attainable after x minutes under the AP Control device. (a) Original tuning that works well; (b) original tuning that works poorly; and (c) results of improved tuning for the patient and bin in (b).

values are calculated using the optimization problem in Figure 8. Details for the insulin on board calculation, and soundness theorems can be found in Kushner et al [91].

7.3.1 Data-Driven Controller Tuning

Beyond just verifying the published approach to selecting the PID controller parameters [150], Kushner et al investigated how to tune the gains of the controller to potentially improve the time in range for the controller.

The tuning rule proposed by Weinzimer et al first selects the proportional gain K_p. Once K_p is set, the value of integral gain is set to $K_i = K_p T_i$ and derivative gain $K_d = K_p / T_d$. Specifically, K_p is chosen as the patient's daily insulin dose divided by 135. However, in practice this rule-of-thumb base tuning is known to result in sub-optimal performance in terms of time in range. In some instances, the reachability analysis observed convergence of all possible behaviors of the closed loop to a desirable range (see Fig. 9(a)), whereas others do not (see Fig. 9(b)). However, in the latter case, adjusting the K_p parameter using the verification results can result in drastically improved outcomes (see 9(c)).

Rather than search for parameters for each patient individually, it is possible to derive a new *tuning law* in terms of a formula that suggests a value of K_p parameter based on the patient's daily insulin requirement or HbA1C readings. In order to identify such an improved tuning law, an exhaustive set of possible tuning parameters were selected for each individual, and the reachability analysis procedure was run for each. After this, the optimal tuning parameter is identified as that which enables the device to maintain the individual's blood glucose both within the range $70-180$mg/dL, and as close to the euglycemic range, $80-120$mg/dL, as possible. Individuals are grouped based on how much re-tuning they require. Analysis was performed to determine whether an individual's demographics are correlated with this proportion of required re-tuning, and a new tuning protocol was identified based on a combination if individual's HbA1c levels as well as daily insulin requirements. Details can be found in Kushner et al [91]. Under this new tuning, time-in-range (time with blood glucose between $70-180$mg/dL) was improved by 83%, and instances of hypo and hyperglycemia were significantly lessened (Table 3).

Table 3 Changes in % time in range: original tuning rule vs. the new tuning rule.

Time Spent in Category	**Original Tuning Rule** mean(std)	**New Tuning Rule** mean(std)
Safe Range 70−180mg/dL	14.80% (21.60%)	84.77% (26.54%)
Severe Hypoglycemia BG< 50mg/dL	29.14% (33.36%)	0% (0%)
Hypoglycemia BG< 70mg/dL	85.12% (21.56%)	15.23% (26.54%)
Hyperglycemia BG> 250mg/dL	0% (0%)	0% (0%)

7.4 Verification of AP: Future Directions

The problem of verifying AP has received some attention with promising initial results. At the same time, formal specification and verification approaches have remained mostly irrelevant to the development of AP control systems. First of all, whereas existing approaches such as S-Taliro are theoretically capable of handling a large class of controllers, they suffer from the problem of "simulation explosion" wherein (a) simulations become more expensive to perform at a higher level of fidelity, and (b) the number of simulations needed to explore all corner cases is prohibitively high as the model is capable of generating richer sets of behaviors. Furthermore, the problem of variability between patients requires sophisticated data driven models. Currently, verification approaches that tackle data driven models are restricted to simple linear models that can be readily encoded into existing constraint solvers such as SAT-modulo theory (SMT) solvers or MILP solvers. The use of neural networks or probabilistic generative models for capturing insulin-glucose response in patients is currently being investigated by many groups including ours, and holds the promise for improved verification results.

Another gap lies in explaining the results of verification to engineers and clinicians. This involves explaining the violations found by the verification in terms of causal factors or root causes. Work by Diwakaran et al demonstrates the analysis of neighborhoods from falsifying traces to measure the influence of various inputs on the violation of a property [44]. However, in practice this approach can be prohibitively expensive.

8 Conclusion

We conclude by looking forward to upcoming new features in the *next generation* artificial pancreas algorithms which will pose new challenges for verification. The next generation AP will be an autonomous system that *completely* controls the insulin delivery, reducing the user burden to the minimum extent feasible. Currently, the main aspects include proactively bolusing for future meals, shutting off insulin delivery in advance of strenuous physical activities, periodically calibrating the CGM using a "finger-stick" glucometer, delivering correction boluses to reduce high BG levels during meals, managing nighttime low BG by setting alarms, and changing insulin delivery during periods of fever, stress, and other physical conditions which affect insulin sensitivity. The next generation devices seek to partially, or totally, eliminate many of these interventions. First, rather than rely

on a single sensor these devices will rely on a variety of available sensors including a "smart watch" that can sense physical activity, heart rate, body position, a GPS that can sense physical location meaningful to BG levels such as a gym, soccer field or a restaurant, specialized sensors such as Google's contact lens BG meter [122], an in-ear sensor that can potentially sense jaw movement [118], and a smart fork that can be used to eat and thus inform the system about incoming meals [45, 72, 100]. At the same time, improvements to natural language understanding have led to systems capable of communicating with their users through natural language speech/text interfaces. The potential for these advances to improve BG control is already being seen in early results [26, 83, 148].

Numerous critical challenges in specification and verification arise from these developments. The use of prediction models allows us to use MPC so that desirable properties can "baked in" by construction, and thus, runtime monitors can be used to detect inconsistencies between the forecasts and reality over time, responding to deviations which can threaten patient well-being. To conclude, we have described the basic pathophysiology of type-1 diabetes, its treatment, the AP project, modeling approaches, control algorithms, and specification and verification methodologies. Rapid advances in biomedicine combined with data-driven AI approaches are revolutionizing this field with the promise of immense benefits to people with type-1 diabetes and insulin-dependent type-2 diabetes. At the same time, numerous challenges present themselves that will continue to engage researchers in automated reasoning and formal methods as a whole for the foreseeable future.

Acknowledgments: The authors gratefully acknowledge detailed comments from the anonymous reviewers. This work was supported in part by the US National Science Foundation (NSF) under grant numbers 1446900, 1446751 and 1646556. All opinions expressed are those of the authors and not necessarily of the NSF.

References

[1] H. Abbas, G. Fainekos, S. Sankaranarayanan, F. Ivancic, and A. Gupta. Probabilistic temporal logic falsification of cyber-physical systems. *Trans. on Embedded Computing Systems (TECS)*, 12:95–, 2013.

[2] R. Advisory. R7-2016-07: Multiple vulnerabilities in animas onetouch ping insulin pump, 2016. Cf. `https://community.rapid7.com/community/infosec/blog/2016/10/04/r7-2016-07-multiple-vulnerabilities-in-animas-onetouch-ping-insulin-pump`.

[3] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

[4] K. Alberti and P. Zimmet. Definition, diagnosis and classification of diabetes mellitus and its complications. part 1: diagnosis and classification of diabetes mellitus. provisional report of a who consultation. *Diabetic Medicine*, 15(7): 539–553, Jul 1998.

[5] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. Molecular biology of the cell, (garland science, new york, 2008). *Google Scholar*, page 652, 2002.

[6] Y. S. R. Annapureddy, C. Liu, G. E. Fainekos, and S. Sankaranarayanan. S-taliro: A tool for temporal logic falsification for hybrid systems. In *Tools and algorithms for the construction and analysis of systems*, volume 6605 of *LNCS*, pages 254–257. Springer, 2011.

[7] E. Atlas, R. Nimri, S. Miller, E. A. Grunberg, and M. Phillip. MD-Logic artificial pancreas system: A pilot study in adults with type 1 diabetes. *Diabetes Care*, 33 (5):1072–1076, May 2010.

[8] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

[9] R. Basu, B. Di Camillo, G. Toffolo, A. Basu, P. Shah, A. Vella, R. Rizza, and C. Cobelli. Use of a novel triple-tracer approach to assess postprandial glucose metabolism. *American Journal of Physiology-Endocrinology And Metabolism*, 284(1):E55–E69, 2003.

[10] N. Baysal, F. Cameron, B. A. Buckingham, D. M. Wilson, H. P. Chase, D. M. Maahs, and B. Bequette. A novel method to detect pressure-induced sensor attenuations (PISA) in an artificial pancreas. *J. Diabetes Science and Technology*, 8(6):1091–1096, 2014.

[11] B. Bequette, F. Cameron, B. Buckingham, D. Maahs, and J. Lum. Overnight hypoglycemia and hyperglycemia mitigation for individuals with type 1 diabetes. how risks can be reduced. *IEEE Control Systems*, pages 125–134, February 2018. doi: 10.1109/MCS.2017.2767119.

[12] B. W. Bequette. Algorithms for a closed-loop artificial pancreas: The case for model predictive control. *J. Diabetes Science and Technology*, 7:1632–1643, November 2013.

[13] R. N. Bergman. Minimal model: Perspective from 2005. *Hormone Research*, pages 8–15, 2005. doi: 10.1159/000089312.

[14] R. N. Bergman. Orchestration of glucose homeostasis: from a small acorn to the california oak. *Diabetes*, 56(6):1489–501, June 2007.

[15] R. N. Bergman and J. Urquhart. The pilot gland approach to the study of insulin secretory dynamics. *Recent Progress in Hormone Research*, 27:583–605, 1971.

[16] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli. Quantitative estimation of insulin sensitivity. *American Journal of Physiology-Endocrinology and Metabolism*, 236(6):E667, Jun 1979.

[17] D. Bertsimas, V. Gupta, and N. Kallus. Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292, 2018.

[18] V. W. Bolie. Coefficients of normal blood glucose regulation. *Journal of Applied Physiology*, 16(5):783–788, Sep 1961.

[19] A. Borri, F. Cacace, A. De Gaetano, A. Germani, C. Manes, P. Palumbo, S. Panunzi, and P. Pepe. Luenberger-like observers for nonlinear time-delay systems with application to the artificial pancreas: The attainment of good performance. *IEEE Control Systems*, 37(4):33–49, 2017.

[20] F. Cameron. *Explicitly Minimizing Clinical Risk through Closed-loop Control of Blood Glucose in Patients with Type 1 Diabetes Mellitus*. PhD thesis, Stanford University, 2010.

[21] F. Cameron, B. W. Bequette, D. Wilson, B. Buckingham, H. Lee, and G. Niemeyer. Closed-loop artificial pancreas based on risk management. *J. Diabetes Sci Technol.*, 5(2):368–79, 2011.

[22] F. Cameron, G. Niemeyer, and B. W. Bequette. Extended multiple model prediction with application to blood glucose regulation. *Journal of Process Control*, 22(8): 1422–1432, Sep 2012.

[23] F. Cameron, D. M. Wilson, B. A. Buckingham, H. Arzumanyan, P. Clinton, H. P. Chase, J. Lum, D. M. Maahs, P. M. Calhoun, and B. W. Bequette. Inpatient studies of a kalman-filter-based predictive pump shutoff algorithm. *J. Diabetes Science and Technology*, 6(5):1142–1147, 2012.

[24] F. Cameron, G. Niemeyer, D. M. Wilson, B. W. Bequette, K. S. Benassi, P. Clinton, and B. A. Buckingham. Inpatient Trial of an Artificial Pancreas Based on Multiple Model Probabilistic Predictive Control with Repeated Large Unannounced Meals. *Diabetes Technology & Therapeutics*, pages 728–734, Sept. 2014. doi: 10.1089/dia.2014.0093.

[25] F. Cameron, G. Fainekos, D. M. Maahs, and S. Sankaranarayanan. Towards a verified artificial pancreas: Challenges and solutions for runtime verification. In *Proceedings of Runtime Verification (RV'15)*, volume 9333 of *Lecture Notes in Computer Science*, pages 3–17, 2015.

[26] F. M. Cameron, T. T. Ly, B. A. Buckingham, D. M. Maahs, G. P. Forlenza, C. J. Levy, D. Lam, P. Clinton, L. H. Messer, E. Westfall, C. Levister, Y. Y. Xie, N. Baysal, D. Howsmon, S. D. Patek, and B. Bw. Closed-loop control without meal announcement in type 1 diabetes. *Diabetes Technol Therapeutics*, 19(9):527–532, 2017. doi: 10.1089/dia.2017.0078.

[27] H. P. Chase and D. Maahs. *Understanding Diabetes (Pink Panther Book)*. Children's Diabetes Foundation, 12 edition, 2011. Available online through CU Denver Barbara Davis Center for Diabetes.

[28] F. Chee and T. Fernando. *Closed-Loop Control of Blood Glucose*. Springer, 2007.

[29] S. Chen, M. O'Kelly, J. Weimer, O. Sokolsky, and I. Lee. An intraoperative glucose control benchmark for formal verification. In 5^{th} *IFAC conference on Analysis and Design of Hybrid Systems (ADHS)*, Oct 2015.

[30] X. Chen, E. Ábrahám, and S. Sankaranarayanan. Flow*: An analyzer for non-linear hybrid systems. In *Proc. of CAV'13*, volume 8044 of *LNCS*, pages 258–263. Springer, 2013.

[31] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 1999.

[32] W. L. Clarke, S. Anderson, M. Breton, S. Patek, L. Kashmer, and B. Kovatchev. Closed-loop artificial pancreas using subcutaneous glucose sensing and insulin delivery and a model predictive control algorithm: The virginia experience. *Journal of Diabetes Science and Technology*, 3(5):1031–1038, Sep 2009.

[33] Claudio Cobelli et al. and AP@Home Consortium. First use of model predictive control in outpatient wearable artificial pancreas. *Diabetes Care*, 37(5):1212–1215, May 2014.

[34] C. Cobelli, D. Foster, and G. Toffolo. *Tracer Kinetics in Biomedical Research*. Springer Science & Business Media, 2000.

[35] C. Cobelli, C. D. Man, G. Sparacino, L. Magni, G. D. Nicolao, and B. P. Kovatchev. Diabetes: Models, signals and control (methodological review). *IEEE reviews in biomedical engineering*, 2:54–95, 2009.

[36] C. Cobelli, E. Renard, and B. Kovatchev. Artificial pancreas: Past, present, future. *Diabetes Care*, 60(11):2672–2682, November 2011.

[37] D. A. Copp, R. Gondhalekar, and J. P. Hespanha. Simultaneous model predictive control and moving horizon estimation for blood glucose regulation in type 1 diabetes. *Optimal Control Applications and Methods*, 39(2):904–918, 2018.

[38] P. E. Cryer. Hypoglycemia, functional brain failure, and brain death. *Journal of Clinical Investigation*, 117(4):868–870, Apr 2007.

[39] C. Cutler and B. Ramaker. Dynamic matrix controlâĂŤa computer control algorithm. In *Proceedings of the joint automatic control conference*, 1980. Paper WP5-B.

[40] C. Dalla Man, M. Camilleri, and C. Cobelli. A system model of oral glucose absorption: validation on gold standard data. *IEEE Trans. on Biomedical Engg.*, 53(12):2472–2478, 2006.

[41] C. Dalla Man, M. D. Breton, and C. Cobelli. Physical activity into the meal glucoseÂŮinsulin model of type 1 diabetes: in silico studies, 2009.

[42] E. Dassau, H. Zisser, R. Harvey, M. Percival, B. Grosman, W. Bevier, E. Atlas, S. Miller, R. Nimri, L. Jovanovic, and F. Doyle. Clinical evaluation of a personalized artificial pancreas. *Diabetes Care*, 36(4):801–9, 2013.

[43] L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.

[44] R. Diwakaran, S. Sankaranarayanan, and A. Trivedi. Analyzing neighbourhoods of falsifying traces, 2017. Intl. Conference on CPS (to appear).

[45] Y. Dong, A. Hoover, J. Scisco, and E. Muth. A new method for measuring meal intake in humans via automated wrist motion tracking. *Applied Psychophysiology and Biofeedback*, 37(3):205–215, 2012.

[46] A. Donzé. Breach: A toolbox for verification and parameter synthesis of hybrid systems. In *CAV*, volume 6174 of *Lecture Notes in Computer Science*. Springer, 2010.

[47] A. Donzé and O. Maler. Robust satisfaction of temporal logic over real-valued signals. In *FORMATS*, volume 6246 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2010.

[48] F. J. Doyle, L. M. Huyett, J. B. Lee, H. C. Zisser, and E. Dassau. Closed-loop artificial pancreas systems: Engineering the algorithms. *Diabetes Care*, 37:1191–1197, 2014.

[49] A. Dunaif and D. T. Finegood. Beta-cell dysfunction independent of obesity and glucose intolerance in the polycystic ovary syndrome. *The Journal of Clinical Endocrinology and Metabolism*, 81(3):942–947, 1996.

[50] S. Dutta, T. Kushner, and S. Sankaranarayanan. Robust data-driven control of artificial pancreas systems using neural networks. In *International Conference on Computational Methods in Systems Biology*, pages 183–202. Springer, 2018.

[51] F. El-Khatib, J. Jiang, and E. R. Damiano. Adaptive closed-loop control provides blood-glucose regulation using dual subcutaneous insulin and glucagon infusion in diabetic swine. *J Diabetes Sci Technol.*, 1(2):181–92, 2007.

[52] F. H. El-Khatib, S. J. Russell, D. M. Nathan, R. G. Sutherlin, and E. R. Damiano. A bihormonal closed-loop artificial panceras for type 1 diabetes. *Sci. Transl. Med.*, 2, April 2010.

[53] A. Facchinetti, G. Sparacino, and C. Cobelli. Modeling the error of continuous glucose monitoring sensor data: Critical aspects discussed through simulation studies. *J. Diabetes Sci. and Tech.*, 4(1), January 2010.

[54] G. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410:4262–4291, 2009.

[55] G. Forlenza, S. Deshpande, T. Ly, D. Howsmon, F. Cameron, N. Baysal, E. Mauritzen, T. Marcal, L. Towers, B. Bequette, L. Huyett, J. Pinsker, R. Gondhalekar, F. I. Doyle, D. Maahs, B. Buckingham, and E. Dassau. Application of zone model predictive control artificial pancreas during extended use of infusion set and sensor: A randomized crossover-controlled home-use trial. *Diabetes Care*, 40:1096–1102, 2017. `https://doi.org/10.2337/dc17-0500`.

[56] G. Forlenza, F. Cameron, T. Ly, D. Lam, D. Howsmon, N. Baysal, G. Kulina, L. Messer, P. Clinton, C. Levister, S. Patek, C. Levy, R. Wadwa, D. Maahs, B. Bequette, and B. B. Fully closed-loop multiple model predictive controller (mmppc) artificial pancreas (ap) performance in adolescents and adults in a supervised hotel setting. *Diabetes Technol. Therapeutics*, 20:5, 2018. doi: 10.1089/dia.2017.0424.

[57] C. Fraley and A. E. Raftery. How many clusters? which clustering method? answers via model-based cluster analysis. *The computer journal*, 41(8):578–588, 1998.

[58] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler. SpaceEx: Scalable verification of hybrid systems. In *Proc. CAV'11*, volume 6806 of *LNCS*, pages 379–395, 2011.

[59] S. Gao, S. Kong, and E. M. Clarke. dReal: an SMT solver for nonlinear theories over the reals. In *Proc. CADE'13*, volume 7898 of *Lecture Notes in Computer Science*, pages 208–214. Springer, 2013.

[60] G. Garcia and A. Morshedi. Quadratic programming solution of dynamic matrix control (QDMC). *Chem Eng. Commun.*, 46:73–87, 1986.

[61] S. K. Garg, S. A. Weinzimer, W. V. Tamborlane, B. A. Buckingham, and *others*. Glucose outcomes with the in-home use of a hybrid closed-loop insulin delivery system in adolescents and adults with type 1 diabetes. *Diabetes Technology and Therapeutics*, 19(3):1–9, 2017.

[62] E. I. Georga, V. C. Protopappas, D. Polyzos, and D. I. Fotiadis. A predictive model of subcutaneous glucose concentration in type 1 diabetes based on random forests. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pages 2889–2892. IEEE, 2012.

[63] E. I. Georga, V. C. Protopappas, D. Ardigò, M. Marina, I. Zavaroni, D. Polyzos, and D. I. Fotiadis. Multivariate prediction of subcutaneous glucose concentration in type 1 diabetes patients based on support vector regression. *IEEE journal of biomedical and health informatics*, 17(1):71–81, 2013.

[64] M. Ghorbani and P. Bogdan. Reducing risk of closed loop control of blood glucose in artificial pancreas using fractional calculus. In *36th Annual International*

Conference of the IEEE Engineering in Medicine and Biology Society (EMBS), pages 4839–4842, Aug 2014.

[65] R. Gondhalekar, E. Dassau, and F. J. Doyle. Moving-horizon-like state estimation via continuous glucose monitor feedback in mpc of an artificial pancreas for type 1 diabetes. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 310–315. IEEE, 2014.

[66] R. Gondhalekar, E. Dassau, and F. J. Doyle III. Periodic zone-mpc with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes. *Automatica*, 71:237–246, 2016.

[67] L. Griva, M. Breton, D. Chernavvsky, and M. Basualdo. Commissioning procedure for predictive control based on arx models of type 1 diabetes mellitus patients. *IFAC-PapersOnLine*, 50(1):11023–11028, 2017.

[68] B. Grosman, E. Dassau, H. Zisser, L. Jovanovic, and F. Doyle. Zone model predictive control: A strategy to minimize hyper- and hypoglycemic events. *J Diabetes Sci Technol.*, 4(4):961–75, 2010.

[69] B. Grosman, E. Dassau, H. C. Zisser, L. Jovanovič, and F. J. Doyle. Zone model predictive control: A strategy to minimize hyper- and hypoglycemic events. *Journal of Diabetes Science and Technology*, 4(4):961–975, Jul 2010.

[70] B. Grosman, D. Wu, D. Miller, L. Lintereur, A. Roy, N. Parikh, and F. R. Kaufman. Sensor-augmented pump-based customized mathematical model for type 1 diabetes. *Diabetes technology & therapeutics*, 20(3):207–221, 2018.

[71] D. M. Hamby. A review of techniques for parameter sensitivity analysis of environmental models. *Environmental Monitoring Assessment*, 32(2):135–154, September 1994.

[72] HAPIfork. HAPIfork. `https://www.hapi.com/product/hapifork`. Accessed: February 26, 2017.

[73] R. Harvey, E. Dassau, and *et al.* Clinical evaluation of an automated artificial pancreas using zone-model predictive control and health monitoring system. *Diabetes Technol. Ther.*, 16:348âĂŞ357, 2014.

[74] Hooman Hakami (Medtronic Inc.). FDA approves MINIMED 670G system - world's first hybrid closed loop system. URL `https://www.medtronicdiabetes.com/blog/fda-approves-minimed-670g-system-worlds-first-hybrid-closed-loop-system/`.

[75] R. Hovorka. Continuous glucose monitoring and closed-loop systems. *Diabetic Medicine*, 23(1):1–12, 2005.

[76] R. Hovorka, F. Shojaee-Moradie, P. Carroll, L. Chassin, I. Gowrie, N. Jackson, R. Tudor, A. Umpleby, and R. Hones. Partitioning glucose distribution/transport, disposal and endogenous production during IVGTT. *Am. J. Physiol. Endocrinol. Metab.*, 282:992–1007, 2002.

[77] R. Hovorka, V. Canonico, L. Chassin, U. Haueter, M. Massi-Benedetti, M. Frederici, T. Pieber, H. Shaller, L. Schaupp, T. Vering, and M. Wilinska. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiological Measurement*, 25:905–920, 2004.

[78] R. Hovorka, J. M. Allen, D. Elleri, L. J. Chassin, J. Harris, D. Xing, C. Kollman, T. Hovorka, A. M. Larsen, M. Nodale, A. D. Palma, M. Wilinska, C. Acerini, and

D. Dunger. Manual closed-loop delivery in children and adoloscents with type 1 diabetes: a phase 2 randomised crossover trial. *Lancet*, 375:743–751, February 2010.

[79] D. P. Howsmon, F. Cameron, N. Baysal, T. T. Ly, G. P. Forlenza, D. M. Maahs, B. A. Buckingham, J. Hahn, and B. W. Bequette. Continuous glucose monitoring enables the detection of losses in infusion set actuation (LISAs). *Sensors*, 2017:17. doi: 10.3390/s17010161.

[80] D. P. Howsmon, N. Baysal, B. A. Buckingham, G. P. Forlenza, T. T. Ly, D. M. Maahs, T. Marcal, L. Towers, E. Mauritzen, S. Deshpande, L. M. Huyett, J. E. Pinsker, R. Gondhalekar, F. J. D. III, E. Dassau, J. Hahn, and B. W. Bequette. Real-time detection of infusion site failures in a closed-loop artificial pancreas. *Diabetes Science and Technology*, 2018. URL `https://doi.org/10.1177/19322968187551`. Online.

[81] F. J. D. III, L. M. Huyett, J. B. Lee, H. C. Zisser, and E. Dassau. Closed-loop artificial pancreas systems: Engineering the algorithms. *Diabetes Care*, 37(5): 1191–1197, May 2014.

[82] W. J, R. R, and B. T. Guidelines for insulin dosing in continuous subcutaneous insulin infusion using new formulas from a retrospective study of individuals with optimal glucose levels. *J Diabetes Sci Technol*, 4:1174–81, 2010.

[83] P. G. Jacobs, N. Resalat, J. El Youssef, R. Reddy, D. Branigan, N. Preiser, J. Condon, and J. Castle. Incorporating an exercise detection, grading, and hormone dosing algorithm into the artificial pancreas using accelerometry and heart rate. *Journal of diabetes science and technology*, 9(6):1175–1184, 2015.

[84] T. Jayalakshmi and A. Santhakumaran. A novel classification method for diagnosis of diabetes mellitus using artificial neural networks. In *Data Storage and Data Engineering (DSDE), 2010 International Conference on*, pages 159–163. IEEE, 2010.

[85] S. M. Kissler, C. Cichowitz, S. Sankaranarayanan, and D. M. Bortz. Determination of personalized diabetes treatment plans using a two-delay model. *J. Theor. Biol.*, (accepted), 2014.

[86] M. T. Korytkowski, S. L. Berga, and M. J. Horwitz. Comparison of the minimal model and the hyperglycemic clamp for measuring insulin sensitivity and acute insulin response to glucose. *Metabolism*, 44(9):1121–1125, 1995.

[87] B. Kovatchev, C. Cobelli, E. Renard, S. Anderson, M. Breton, S. Patek, W. Clarke, D. Bruttomesso, A. Maran, S. Costa, A. Avogaro, C. Dalla Man, A. Facchinetti, L. Magni, G. De Nicolao, J. Place, and A. Farret. Multinational study of subcutaneous model-predictive closed-loop control in type 1 diabetes mellitus: summary of the results. *J Diabetes Sci Technol.*, 4(6):1374–81, 2010.

[88] B. P. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli. In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes, 2009.

[89] A. Kowalski. Pathway to artificial pancreas revisited: Moving downstream. *Diabetes Care*, 38:1036–1043, June 2015.

[90] R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.

[91] T. Kushner, D. Bortz, D. Maahs, and S. Sankaranarayanan. A data-driven approach to artificial pancreas verification and synthesis. In *Intl. Conference on Cyber-Physical Systems (ICCPS'18)*. IEEE Press, 2018.

[92] J. Kusunoki, A. Kanatani, and D. E. Moller. Modulation of fatty acid metabolism as a potential approach to the treatment of obesity and the metabolic syndrome. *Endocrine*, 29(1):91–100, Feb 2006.

[93] H. Lee and B. Bequette. A closed-loop artificial pancreas based on MPC: human-friendly identification and automatic meal disturbance rejection. *Biomedical Signal Processing and Control*, 4(4):347–354, 2009.

[94] H. Lee, B. Buckingham, D. Wilson, and B. Bequette. A closed-loop artificial pancreas using model predictive control and a sliding meal size estimator. *J. Diabetes Sci. Tech.*, 3(5):1082–1090, 2009.

[95] E. Lehmann and T. Deutsch. A physiological model of glucose-insulin interaction in type 1 diabetes mellitus. *Journal of biomedical engineering*, 14(3):235–242, 1992.

[96] C. Li, A. Raghunathan, and N. K. Jha. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. In *Intl. Conf. on e-Health Networking, Applications and Security*, pages 151–156, 2011.

[97] J. Li, Y. Kuang, and B. Li. Analysis of ivgtt glucose-insulin interaction models with time delay. *Discrete and Continuous Dynamical Systems Series B*, 1(1):103–124, 2001.

[98] J. Li, Y. Kuang, and C. C. Mason. Modeling the glucose-insulin regulatory system and ultradian insulin secretory oscillations with two explicit time delays. *J. Theor. Biol.*, 242(3):722–35, Oct. 2006.

[99] J. Liepe, P. Kirk, S. Filippi, T. Toni, C. P. Barnes, and M. P. H. Stumpf. A framework for parameter estimation and model selection from experimental data in systems biology using approximate bayesian computation. *Nature Protocols*, 9(2):439–456, January 2014.

[100] J. Liu, E. Johns, L. Atallah, C. Pettitt, B. Lo, G. Frost, and G. Z. Yang. An intelligent food-intake monitoring system using wearable sensors. In *2012 Ninth International Conference on Wearable and Implantable Body Sensor Networks*, pages 154–160, 2012.

[101] K. Lunze, T. Singh, M. Walter, M. D. Brendel, and S. Leonhardt. Blood glucose control algorithms for type 1 diabetic patients: A methodological review. *Biomedical Signal Processing and Control*, 8(2):107 – 119, 2013. ISSN 1746-8094.

[102] D. M. Maahs, P. Calhoun, B. A. Buckingham, and Others. A randomized trial of a home system to reduce nocturnal hypoglycemia in type 1 diabetes. *Diabetes Care*, 37(7):1885–1891, July 2014.

[103] L. Magni, D. Raimondo, L. Bossi, C. D. Man, G. D. Nicolao, B. Kovatchev, and C. Cobelli. Model predictive control of type 1 diabetes: an *in silico* trial. *J. Diabetes Science and Technology*, 1(6):804–12, 2007.

[104] Z. Mahmoudi, F. Cameron, N. K. Poulsen, H. Madsen, B. W. Bequette, and J. B. Jørgensen. Sensor-based detection and estimation of meal carbohydrates for people with diabetes. *Biomedical Signal Processing and Control*, 48:12–25, 2019.

[105] A. Makroglou, J. Li, and Y. Kuang. Mathematical models and software tools for the glucose-insulin regulatory system and diabetes: an overview. *Applied numerical mathematics*, 56(3-4):559–573, 2006.

[106] O. Maler and D. Nickovic. Monitoring Temporal Properties of Continuous Signals. In *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, pages 152 – 166. Springer, 2004.

[107] C. Man, M. Camilleri, and C. Cobelli. A system model of oral glucose absorption: Validation on gold standard data. *Biomedical Engineering, IEEE Transactions on*, 53(12):2472 –2478, dec. 2006.

[108] C. D. Man, R. A. Rizza, and C. Cobelli. Meal simulation model of the glucose-insulin system. *IEEE Transactions on Biomedical Engineering*, 1(10):1740–1749, 2006.

[109] C. D. Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli. The uva/padova type 1 diabetes simulator: new features. *Journal of diabetes science and technology*, 8(1):26–34, 2014.

[110] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer, New York, 1995.

[111] G. Marchetti, M. Barolo, L. Jovanovič, H. Zisser, and D. E. Seborg. A feedforward–feedback glucose control strategy for type 1 diabetes mellitus. *Journal of Process Control*, 18(2):149 – 162, 2008.

[112] E. Marieb and K. Hoehn. Human anatomy and physiology 2004. *San Francisco, Daryl Fox*.

[113] R. Mauseth, Y. Wang, E. Dassau, R. Kircher, D. Matheson, H. Zisser, and *others*. Proposed clinical application for tuning fuzzy logic controller of artificial pancreas utilizing a personalization factor. *J Diabetes Sci Technol*, (4):913–922, 2010.

[114] Medtronic Inc. "paradigm" insulin pump with low glucose suspend system, 2012. Cf. `http://www.medtronicdiabetes.ca/en/paradigm_veo_glucose.html`.

[115] N. Musi and L. J. Goodyear. Insulin resistance and improvements in signal transduction. *Endocrine*, 29(1):73–80, Feb 2006.

[116] K. R. Muske and T. A. Badgwell. Disturbance modeling for offset-free linear model predictive control. *J Process Control*, 12:617–32, 2002.

[117] T. Nghiem, S. Sankaranarayanan, G. E. Fainekos, F. Ivančić, A. Gupta, and G. J. Pappas. Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In *Hybrid Systems: Computation and Control*, pages 211–220. ACM Press, 2010.

[118] A. Nguyen, R. Alqurashi, Z. Raghebi, F. Banaei-kashani, A. C. Halbower, and T. Vu. A lightweight and inexpensive in-ear sensing system for automatic whole-night sleep stage monitoring. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, SenSys '16, pages 230–244, 2016.

[119] G. D. Nicolao, L. Magni, C. D. Man, and C. Cobelli. Modeling and control of diabetes: Towards the artificial pancreas. *IFAC Proceedings Volumes*, 44(1):7092 – 7101, 2011. 18th IFAC World Congress.

[120] R. Nimri, I. Muller, E. Atlas, S. Miller, O. Kordonouri, N. Bratina, C. Tsioli, M. Stefanija, T. Danne, T. Battelino, and P. M. Night glucose control with md-logic

artificial pancreas in home setting: a single blind, randomized crossover trial-interim analysis. *Pediatric Diabetes*, 15(2):91–100, March 2014.

[121] G. Nucci and C. Cobelli. Models of subcutaneous insulin kinetics. a critical review. *Computer methods and programs in biomedicine*, 62(3):249–257, 2000.

[122] B. Otis and B. Parviz. Introducing google's smart contact lens project. Blog post on Google Inc. official weblog, `http://googleblog.blogspot.com/2014/01/introducing-our-smart-contact-lens.html`.

[123] N. Paoletti, K. S. Liu, S. A. Smolka, and S. Lin. Data-driven robust control for type 1 diabetes under meal and exercise uncertainties. In *Computational Methods in Systems Biology (CMSB)*, volume 10545 of *Lecture Notes in Computer Science*, pages 214–232. Springer, 2017.

[124] R. S. Parker, F. J. Doyle III, J. H. Ward, and N. A. Peppas. Robust h glucose control in diabetes using a physiological model. *AIChE Journal*, 46(12):2537–2549, 2000.

[125] R. S. Parker, F. J. Doyle, and N. A. Peppas. The intravenous route to blood glucose control. *IEEE Engineering in Medicine and Biology Magazine*, 20(1):65–73, 2001.

[126] S. Patek, B. Bequette, M. Breton, B. Buckingham, E. Dassau, F. Doyle III, J. Lum, L. Magni, and H. Zisser. In silico preclinical trials: methodology and engineering guide to closed-loop control in type 1 diabetes mellitus. *J Diabetes Sci Technol.*, 3 (2):269–82, 2009.

[127] C. Pérez-Gandía, A. Facchinetti, G. Sparacino, C. Cobelli, E. Gómez, M. Rigla, A. de Leiva, and M. Hernando. Artificial neural network algorithm for online glucose prediction from continuous glucose monitoring. *Diabetes technology & therapeutics*, 12(1):81–88, 2010.

[128] G. Pillonetto, G. Sparacino, and C. Cobelli. Numerical non-identifiability regions of the minimal model of glucose kinetics: superiority of bayesian estimation. *Mathematical biosciences*, 184(1):53–67, 2003.

[129] J. E. Pinsker, J. B. Lee, E. Dassau, D. E. Seborg, P. K. Bradley, R. Gondhalekar, W. C. Bevier, L. Huyett, H. C. Zisser, and F. J. Doyle. Randomized crossover comparison of personalized mpc and pid control algorithms for the artificial pancreas. *Diabetes Care*, 39(7):1135âĂŞ1142, 2016.

[130] A. Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reasoning*, 41 (2):143–189, 2008.

[131] K. Plis, R. C. Bunescu, C. Marling, J. Shubrook, and F. Schwartz. A machine learning approach to predicting blood glucose levels for diabetes management. In *AAAI Workshop: Modern Artificial Intelligence for Health Analytics*, number 31, pages 35–39, 2014.

[132] K. S. Polonsky, J. Sturis, and E. Van Cauter. Temporal profiles and clinical significance of pulsatile insulin secretion. *Hormone Research in Paediatrics*, 49 (3-4):178–184, 1998.

[133] J. Radcliffe. Hacking medical devices for fun and insulin: Breaking the human SCADA system, 2011. Black Hat 2011, Cf.`https://media.blackhat.com/bh-us-11/Radcliffe/BH_US_11_Radcliffe_Hacking_Medical_Devices_WP.pdf`.

[134] C. Ramkissoon, B. Aufderheide, B. W. Bequette, and J. Vehi. Safety and hazards associated with the artificial pancreas. *IEEE Rev. Biomed. Eng.*, 10:44–52, 2017.

[135] J. Rawlings, D. Mayne, and M. Diehl. *Model Predictive Control: Theory, Computation and Design*. Nob Hill Publishing, 2017.

[136] N. Resalat, J. El Youssef, R. Reddy, and P. G. Jacobs. Design of a dual-hormone model predictive control for artificial pancreas with exercise model. In *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*, pages 2270–2273. IEEE, 2016.

[137] J. L. Ruiz, J. L. Sherr, E. Cengiz, L. Carria, A. Roy, G. Voskanyan, W. V. Tamborlane, and S. A. Weinzimer. Effect of insulin feedback on closed-loop glucose control: a crossover study. *Journal of Diabetes Science and Technology*, 6(5):1123–1130, 2012.

[138] M. F. Saad, K. Rebrin, G. M. Steil, et al. Modeling glucose profiles obtained using closed loop insulin delivery—implications for controller optimization. *Diabetes*, 55:A98, 2006.

[139] S. Sankaranarayanan, S. A. Kumar, F. Cameron, B. W. Bequette, G. Fainekos, and D. M. Maahs. Model-based falsification of an artificial pancreas control system. *ACM SIGBED Review (Special Issue on Medical Cyber Physical Systems)*, 2017.

[140] F. Shmarov, N. Paoletti, E. Bartocci, S. Lin, S. Smolka, and P. Zuliani. SMT-based Synthesis of Safe and Robust PID Controllers for Stochastic Hybrid Systems. In *Hardware and Software: Verification and Testing - 13th International Haifa Verification Conference*, pages 131–146. Springer, 2017. doi: 10.1007/978-3-319-70389-3_9. URL `https://link.springer.com/chapter/10.1007%2F978-3-319-70389-3_9`.

[141] M. J. Siper. *An Introduction to Mathematical Theory of Computation (2nd Edition)*. Thompson Publishing (Course Technology), 2005.

[142] J. S. Skyler (editor). *Atlas of Diabetes: Fourth Edition*. Springer Science+Business Media, 2012.

[143] T. Spaic, M. Driscoll, D. Raghiaru, B. Buckingham, D. Wilson, P. Clinton, H. P. Chase, D. Maahs, G. Forlenza, E. Jost, I. Hramiak, T. Paul, B. Bequette, F. Cameron, R. Beck, C. Kollan, J. Lum, and T. Ly. Predictive hyperglycemia and hypoglycemia minimization: In-home evaluation of safety, feasibility, and efficacy in overnight control in type 1 diabetes. *Diabetes Care*, 40(3):359–366, 2017. doi: \url{https://doi.org/10.2337/dc16-1794}.

[144] R. Srinivasan, A. H. Kadish, and R. Sridhar. A mathematical model for the control mechanism of free fatty acid-glucose metabolism in normal humans. *Computers and Biomedical Research*, 3(2):146 – 165, 1970.

[145] G. Steil, A. Panteleon, and K. Rebrin. Closed-sloop insulin delivery - the path to physiological glucose control. *Advanced Drug Delivery Reviews*, 56(2):125–144, 2004.

[146] G. M. Steil. Algorithms for a closed-loop artificial pancreas: The case for proportional-integral-derivative control. *J. Diabetes Sci. Technol.*, 7:1621–1631, November 2013.

[147] K. Turksoy and A. Cinar. Multi-module multivariable artificial pancreas for patients with type 1 diabetes. *IEEE Control Systems Magazine*, 38(1):105–124, 2018.

[148] K. Turksoy, E. S. Bayrak, L. Quinn, E. Littlejohn, and A. Cinar. Multivariable adaptive closed-loop control of an artificial pancreas without meal and activity announcement. *J Diabetes Technology & Therapeutics*, 15(5):386–400, 2013.

[149] K. Turksoy, I. Hajizadeh, S. Samadi, J. Feng, M. Sevil, M. Park, L. Quinn, E. Littlejohn, and A. Cinar. Real-time insulin bolusing for unannounced meals with artificial pancreas. *Control Engineering Practice*, 59:159–164, 2017. doi: 10.1016/j.conengprac.2016.08.001.

[150] S. Weinzimer, G. Steil, K. Swan, J. Dziura, N. Kurtz, and W. Tamborlane. Fully automated closed-loop insulin delivery versus semiautomated hybrid control in pediatric patients with type 1 diabetes using an artificial pancreas. *Diabetes Care*, 31:934–939, 2008.

[151] M. Wilinska, L. Chassin, C. L. Acerini, J. M. Allen, D. Dunber, and R. Hovorka. Simulation environment to evaluate closed-loop insulin delivery systems in type 1 diabetes. *J. Diabetes Science and Technology*, 4, January 2010.

[152] S. Zavitsanou, A. Chakrabarty, E. Dassau, and F. J. Doyle. Embedded control in wearable medical devices: Application to the artificial pancreas. *Processes*, 4(4), 2016.