

Published in final edited form as:

Bioelectron Med (Lond). 2019 March ; 2(1): 29–41. doi:10.2217/bem-2019-0008.

Adding wisdom to ‘smart’ bioelectronic systems: a design framework for physiologic control including practical examples

Aysegul Gunduz¹, Enrico Opri¹, Ro’ee Gilron², Vaclav Kremen³, Gregory Worrell³, Phil Starr², Kent Leyde⁴, Timothy Denison^{*,5}

¹Department of Biomedical Engineering, University of Florida Gainesville, Gainesville, FL 32611, USA

²School of Medicine, University of California San Francisco, San Francisco CA 94143, USA

³Department of Neurology, Mayo Clinic, Rochester, MN 55905, USA

⁴Cadence Neuroscience Inc, Sammamish, WA 98074, USA

⁵Department of Engineering Science, University of Oxford, Oxford, OX3 7DQ, UK

Abstract

This perspective provides an overview of how risk can be effectively considered in physiological control loops that strive for semi-to-fully automated operation. The perspective first introduces the motivation, user needs and framework for the design of a physiological closed-loop controller. Then, we discuss specific risk areas and use examples from historical medical devices to illustrate the key concepts. Finally, we provide a design overview of an adaptive bidirectional brain–machine interface, currently undergoing human clinical studies, to synthesize the design principles in an exemplar application.

Keywords

bioengineering; closed-loop devices; electroceuticals; neuromodulation; neuroprosthesis; smart implants

The advent of ‘smart’ systems technology is revolutionizing many facets of everyday life – from self-driving cars to digital home assistants. However, experiences from this nascent

This work is licensed under the Attribution-NonCommercial-NoDerivatives 4.0 Unported License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>

*Author for correspondence: Tel.: +44 (0)1865 617707; Fax: +44 (0)1865 617728; timothy.denison@eng.ox.ac.uk.

Disclosure

Web address for the corresponding author, T Denison: www.ibme.ox.ac.uk

Financial & competing interests disclosure

This work was supported in part by the NSF 1553482, NIH R01NS096008, and NIH UH3NS095553 for A Gunduz, NIH UH3NS100544 and NIH R01NS090913 for P Starr, NIH R01NS092882 and NIH UH2/UH3NS95495 for G Worrell, and the John Fell Fund and Royal Academy of Engineering (CiET) for T Denison. T Denison, P Starr, G Worrell and K Leyde have IP in this the area of closed loop neurotechnology. G Worrell and Mayo Clinic have a financial interest related to technologies licensed to Cadence Neuroscience Inc. and NeuroOne Inc. K Leyde is CEO of Cadence Neuroscience. The authors have no other relevant affiliations or financial involvement with any organization or entity with a financial interest in or financial conflict with the subject matter or materials discussed in the manuscript apart from those disclosed.

No writing assistance was utilized in the production of this manuscript.

field also provide lessons on the need for the inclusion of risk management in the core system design. Lessons range from the innocuous, such as algorithms exhibiting odd behavior like spontaneously laughing digital assistants, to the severe such as automotive crashes resulting from compromised algorithms, or overconfidence in the automated system's capabilities. As the medical field adapts similar automated technologies [1–3], it is worth considering how engineers and clinical scientists can work together to mitigate risks with deploying intelligent systems through thoughtful system design.

The motivation for 'smart' systems

The emerging field of bioelectronic medicines is diverse, with many examples that motivate intelligent systems that respond more appropriately to physiological signals. These systems range from established cardiac pacemakers with remote monitoring and emergent artificial pancreases that aim to mimic beta cell dynamics, to nascent brain stimulators for epilepsy with adaptive or closed-loop control. Here we use the term adaptive control to refer to systems that respond to a sensory signal, but do not have an explicit feedback loop such as motion-adaptive stimulation [4], while closed-loop control we define as a system with a desired set-point and a response loop that tries to maintain this set-point such as an artificial pancreas [5]. While sometimes used interchangeably, the difference can impact the risk profile, and we attempt to apply the nomenclature carefully. For those interested in more formal reviews of automaticity levels, we refer you to papers such as [6]. In spite of the application diversity and nuances in implementation, the high-level aims are similar for many smart systems. These aims include applying technology such as sensors, machine learning and modeling to efficiently find an optimal therapy or to improve therapy efficacy and reduce side effects, all while striving to lower the burden on the clinical practitioner and the patient consumer.

Deep brain stimulation (DBS) emerged in the 1990s as a neurosurgical intervention for movement disorders [7]. The principle behind DBS is to implant electrodes into deep brain structures such as the basal ganglia and thalamus to modulate the pathological activity causing disabling symptoms through application of electrical pulses (charge or current). The clinical personnel that perform programming of stimulation settings (amplitude, frequency and pulse width of the electrical current) base their decisions on guidelines and the observable behavioral responses and verbal response of patients [8].

While effective strategies exist for selection of stimulation parameters [8], the parameter space is large and it is time consuming and clinically infeasible to evaluate each of the thousands of individual stimulation parameter combinations that may be useful, or optimal, to a given patient. The survey approach can also be problematic in disease states such as epilepsy, where the episodic nature of symptoms of the disease makes an immediate observation of the therapeutic effect of electrical stimulation unlikely. This is further complicated by the fact that self-reported seizure diaries from patients are often inaccurate for consciousness impairing seizures. In effect, the physician is 'flying blind' when adjusting therapy. Even for established therapies like tremor, the introduction of new segmented electrodes further increases programming complexity, placing a greater burden on the clinician. In addition, fixed tonic stimulation, often called 'open loop' stimulation, assumes

that the therapy settings derived in the acute clinical setting serve the patient well for the variations seen outside of clinic; variations in pharmacology, sleep states and other patterns are difficult to capture. Due to these issues, the therapeutic benefit currently achievable with DBS can be strongly dependent on the intuitive skill and experience of the clinician performing the stimulation parameter selection [9].

These challenges motivate technical solutions that can automatically optimize DBS parameters to match the underlying pathology. Studying the neurophysiological signatures of neurological disorders, and the effects of brain stimulation, would enable objective monitoring of fluctuations in the severity of the disorder, and treatment options that can be tailored to the current clinical condition of the patient. A ‘smart’ DBS system could theoretically initiate stimulation and/or adapt its stimulation parameters based on information extracted from brain signals. The guiding hypothesis is that a DBS system capable of detecting the pathological activity, and providing stimulation in an adaptive manner, would provide improved symptom suppression, reduce adverse effects of continuous stimulation and prolong battery life. This in turn could significantly improve the lives of those suffering from neurological disorders.

However, adding automaticity to an implantable brain stimulator can create new issues for delivering therapy. Potential hazards include under- or overstimulation due to a compromised system that puts the patient at risk. To help mitigate these risks systematically, we adopt a framework from regulations originally developed for the ‘physiologic controllers’. Physiologic controllers are defined as ‘devices that try to control a physiological parameter through feedback’. An intuitive example of such a system is an infant incubator where the physiologic control system regulates body temperature using a sensor placed on the infant and adjusts the incubator conditions accordingly. While the standards are prescriptive on where they must be applied, we believe that the intuition gained from framing a smart system in these terms are useful for a broad range of medical device designs.

Frameworks to guide ‘smart’ system design

The design of an ‘intelligent’ implant requires thoughtful consideration for the characteristics of the integrated bioelectronic–physiological system. These considerations are well-captured by a series of technical standards for the safety and essential performance of medical electrical equipment published by the International Electrotechnical Commission (IEC): IEC 60601, 1–10, ‘General requirements for basic safety and essential performance’ – collateral standard: requirements for the development of physiologic closed loop controllers [10]. Using this guidance as a set of design best practices for adaptive and closed-loop responsive medical equipment can help ensure robust operation of a broad class of bioelectronic systems.

Referencing Figure 1, the base bioelectronic system can be modeled as a physiological closed-loop controller with several key subcomponents. The specific terminology derives from the IEC standard, and provides a common language for assessing a system. The core functional blocks of the system are highlighted in blue.

Walking through these block definitions, the patient transfer element is the physiological system that the device is trying to modulate, such as a discrete network of the nervous system. The patient transfer element also includes the disease-specific considerations. For example, the risks associated with the treatment of essential tremor are quite different than those of a cardiac defibrillator. A designer needs to consider the nuances of the particular condition, and their relative impact on technology-derived risks. The actuator is the transduction mechanism by which the bioelectronic system can influence the physiologic system. Examples include electrical stimulation, ultrasound stimulation, transcranial magnetic stimulation or pharmacological agents, including anesthesia. The measurement transfer element is the sensing transducer for detecting a physiologic variable of interest. Examples include bioelectric amplifiers for measuring local field potentials or action potentials, skin impedance, inertial signals, saturated oxygen or interstitial glucose. The command transfer element is the mapping between the desired clinical state and a corresponding variable in the algorithm, translated into the same units as the measuring transfer element and related to where the measurements can be taken. Examples include mapping core body temperature to a target temperature, glucose levels to a desired safe window or a minimum heart rate in a cardiac pacemaker. The comparison element is the block that compares the measured signal against the desired value from the command transfer element. From a control theory perspective, this element estimates the ‘error’ term for feedback mechanisms. The control transfer element is the block that appropriately adjusts the actuator based on the difference between the measured response and the command set-point and can include feedback compensation and similar engineering methods to try and maximize the capability of the overall feedback loop. The patient disturbances capture potential mechanisms that might indirectly influence the control system performance. Typical examples can include drug interactions, circadian rhythms and physical exertion. Similar to the patient transfer element, disturbances usually require disease-specific considerations in the analysis.

Building intuition: mapping historical examples of risk mitigation to the framework

Using the framework as a guide, we can consider: how a physiologic-based control system can introduce new risks, and the actions the designer can take to help mitigate these risks. ‘Risk’, when used in the context of medical device design and regulations is formally defined by the International Standards Organization (ISO) for the application of risk management to medical devices (ISO 14971 ‘Application of risk management to medical devices’). Risk is defined to be a ‘combination of the probability of occurrence of harm and the severity of that harm’; harm is further defined as a ‘physical injury or damage to the health of people, or damage to property or the environment’. Note that the ‘physical’ harm might present itself in a manner that is difficult to observe. For example, mood side effects can arise for brain stimulation in obsessive compulsive disorder which might be difficult to directly monitor with existing sensor approaches [11]. Risk analysis can be used to provide a comprehensive accounting of risk, and as part of an overall risk management process to ensure that risk mitigations are implemented such that risks are reduced to acceptable levels.

Drawing again from the IEC 60601-1-10 framework, additional blocks focused on risks and mitigations are highlighted in brown in Figure 2. We illustrate the core concepts underlying risk and medication using existing medical devices.

First, we consider user-centric intuition for the system, and the need to provide a mental model of the control algorithm. The purpose of the mental model of algorithm operation is to help the practitioner intuitively set up the physiologic controller's parameters, test its operation and resolve issues. For example, a rate response, feed-forward pacemaker has a monotonic mapping of pacing rate to activity levels derived from an embedded motion sensor, with a clinician-defined lower and upper pacing rate limit configured in accordance to the 'actuation limit' consideration [12,13]. The mental model is relatively straightforward: as the patient moves more, you expect the pacing rate for driving their heart should increase to provide the needed cardiac output. The inertial sensor in the device (the measurement transfer element) detects the patient's motion, the comparison element then maps the level of motion to a desired stimulation state. If a clinician wishes to test or configure the system, they can ask the patient to change activity and confirm that the pacing level modulates as expected. Mental models can also have a physiological basis and seek to directly mimic or replace biological functions. For example, control algorithms for the artificial pancreas use dynamic models of pancreatic β cells to help improve their performance [5,14,15]. The clinician can test the performance of this system with challenges (food consumption, activity), and confirm that the insulin delivery provides the appropriate glucose dynamics. Careful consideration of the mental model can help to clearly define key risks associated with the physiologic control loop by: specifying the intended uses cases for the system, defining the human factors of those use cases and identifying the potential hazards.

For risks associated the measuring transfer element, we must focus on how the sensors might cause errors in estimating the patient's state, and the impact of these errors. For example, the designer should understand the expected intra- and interpatient variability. Disease-specific considerations should be included with population characterization to capture the expected breadth of input signals. The goal is to ensure the physiologic loop's measurement has the required dynamic range, signal-to-noise and specificity for the intended patient population. For example, in a cardiac pacemaker or brain interface, this might include the variation in frequency content of the electrocardiogram or neural field potentials, and their signal power (root-mean-square in relevant frequency bands). The variability of signals might require parametric tuning of the signal processing chain and detection classifiers on a disease- and patient-specific basis. Another consideration is detecting when a sensor might be faulty and the measurement transfer element compromised. Mitigations for this can include an independent self-test of the sensor, such as those common in micromachined accelerometers [16], calibration of the sensor [17], redundant sensor systems that provide robustness to a single fault [18–20] or fallback to established safe stimulation parameters when potentially adverse conditions such as electrical interference are detected. Ideally, these mitigations are hidden from the practitioner, and do not add to their burden, which might negatively impact compliance. However, it is critical to provide a system over-ride for patients and physicians where stimulation can be disabled or switched to the predetermined 'safe operating' parameters.

When considering the patient transfer element, the aim is to understand how the physiologic control loop interacts with the patient's physiology and the risks that are associated with biophysical interactions. In practice, the transfer elements require a disease-specific analysis. One example is the impact of response latency on the controller's performance. Latency issues can come from many sources: the temporal dynamics between when an actuation adjustment is made, when this actuation adjustment impacts the physiological system, and when the response might be measured by the sensors. Characterizing and designing for these delays is important, as excess controller bandwidth or sensitivity (gain) can result in undesirable feedback dynamics, such as physiology-controller oscillations. For example, in an artificial pancreas, several minutes might pass before the bolus of insulin takes effect in the body, and several more minutes before the change in blood sugar is perceptible in a subcutaneous glucose sensor placed in the interstitial layer [21,22]. These dynamics can set a limit on the capacity of the system to correct transient perturbations, and motivate feedforward estimates in the model. Latency considerations also lead to considerations for sensitivity and specificity requirements, which are also dependent on which disease state is being treated. More specifically, the disease state will guide the relative penalty for type I (false positive) versus type II (false negative) errors. For example, the fast response timing for a cardiac defibrillator is crucial to maintain blood oxygenation. A fast response time might also suggest a need for high specificity and low latency. These must be balanced against the patient impact of false detections and excessive shocks [23,24]. Given the significance of a missed defibrillation shock, the clinician would probably bias an algorithm to avoid a type II error for a defibrillation at the expense of type I errors. But even type I errors must be kept to an acceptable level for patient acceptance. A similar latency challenge exists for responsive neurostimulators for epilepsy, where the designer wishes to intervene as soon as possible to abort a seizure. The lower impact of false-positive (type I) errors results in a bias toward higher sensitivity, and lower specificity, to minimize the latency to an intervening intervention compared with a cardiac defibrillator [25–27]. Interestingly, the use of hypersensitive seizure detectors with associated marked increase in false-positive seizure electrical stimulation has proven useful and are not clinically penalized since the stimulation is subthreshold, in other words, below patient perception [20–22]. An example where a type I error can result in a negative outcome can be found in epilepsy, where electromagnetic interference or other physiological artifacts might result in excess stimulation for the patient and more rapidly deplete the battery [28]. In sum, the bioelectronic designer must trade-off several considerations for sensitivity, specificity and latency for their specific application – one design does not necessarily serve all needs [1].

Given its direct control of the actuator, there are several risk considerations for the control transfer element. For example, it is essential to define an actuation limit to contain the impact of either excessive or marginal stimulation. More specifically, the designer should consider the impact of actuation levels on the disease-specific physiology of the patient transfer element; note that these can include both higher and lower limits. Returning to the artificial pancreas example, the clinician might want to limit the maximum flow rate of insulin into the body to avoid an overdose, or limit the cumulative insulin dose given over a specified window [5,14,29,30]. As an example of a lower limit, a cardiac pacemaker might

enforce a minimum pacing rate to ensure a patient will always have the minimal cardiac output required to maintain consciousness [12].

Another consideration for the control transfer element is the definition of a fallback ‘safe’ mode. Despite a designer’s best efforts, one of the physiologic controller’s core elements might encounter an issue that undermines performance. This can arise when the measurement transfer element is compromised in some manner. An example of this would be excessive electromagnetic interference saturating a bioelectronic amplifier and thereby leading to misleading information. EMI has in the past been a problem with cardiac pacemakers due to interactions with security wands, microwaves and muscle artifact, and more recently includes the impact of MRI scanners [31–35]. The role of the fallback mode is to allow the physiologic controller to enter a well-defined state with a known safety profile. For example, in the case of a cardiac pacemaker, a fallback mode would be a predetermined pacing rate deemed acceptable to the patient for most situations. This situation might occur when the amplifier’s detection floor rises above a threshold indicating cardiac classification is likely compromised [36].

A final consideration for the control transfer element is a mechanism for providing state indicators. The goal of the state indicator is to provide critical information on the operating mode of the control loop. Given the potential diversity of users, the presentation of this information should include human factors considerations, so that the information is meaningful and actionable. These considerations might also be disease-state dependent. For example, in a typical neurostimulator, the parameter settings for stimulation and state of stimulation are available to both the clinician and the patient, with the data shared being dependent on the specified user. While in other use cases, such as cardiac pacemakers, the patients generally do not have a mechanism to interact with the implant parameters directly. To provide state information, cardiac devices could historically communicate via an embedded auditory signal [37], which the patient could feel, that signaled the need to call a clinician. The limited specificity of this approach could still warrant periodic check-ins [37]. Modern systems are moving to wireless telemetry and remote updates to provide this information directly to the clinician, who then reviews and interprets the diagnostic data and engages the patient as needed [38,39]. In particular, this can be very useful in mitigating risks caused by different system faults, such as an electrode break detected by automated, periodic impedance checks that are logged and if aberrant, passed to the physician via remote communication [40]. In spite of its benefits, a potential hazard of high-level automaticity in state indicators is complacency; designers need to consider the human factors of error checking in their system validation process [41].

A broader consideration for transfer elements, modes and limits is the use of software in modern systems. IEC 62304: ‘Medical device software – software life cycle processes’ establishes that, when considering hazards that can arise from misoperation of a given software element, the probability that the software will eventually misoperate is 100%. For this reason, it is often necessary to implement portions of the control transfer element, fallback mode and actuation limit in hardware, or in separate, well-segregated software elements. Hardware elements may be capable of providing fallback mode behavior independent of software elements. Actuation limits may be implemented as a characteristic

limit of the hardware. Software elements may also be implemented on separate processors. Software elements may be segregated through use of an operating system that provides features such as memory management, guaranteed CPU access and user-space drivers that can be restarted if faults are detected. For example, lock-step processor architectures can be used to provide detection of CPU faults in order to trigger entry into a fallback mode.

Finally, the designer should consider how to understand user error, what causes these errors and ultimately try to avoid them. These aims can be achieved by providing mechanisms for training and fault analysis. Typical design approaches include providing data logs on device usage for training and algorithm ‘debug’. The inclusion of a patient- or algorithm-enabled data recorder can help to understand either how the physiologic control algorithm performs in real-world environments for training purposes, or capture failure events that require modification of the algorithm or additional patient warning. For example, modern cardiac systems store several detected events including the physiological signals, classifier states and actuation settings for assessing field performance [38,39]. These recordings also create a database for continuous algorithm refinement. Additional data logs can include complementary signal sources that assist with understanding the patient environment. For example, the Neurovista seizure prediction system included a microphone on the patient controller that would turn on when the system detected a seizure; this additional information helped to separate seizures with clinical manifestations from those without [42]. These mechanisms have similarities to the ‘black box’ recorders used on airplanes, which help to analyze scenarios that led to a crash and apply this new understanding to either fix a design error or provide additional training to avoid a similar scenario in the future.

An integrated ‘smart’ system case study: bidirectional brain–computer interfacing for treating essential tremor

To pull together the key design elements of a robust physiologic control algorithm, we present a complete system example that illustrates an integrated approach to risk management and mitigations using the IEC 60601, 1–10 based framework we described above. The exemplar is an adaptive deep brain stimulator for the treatment of essential tremor [43], which is one of a family of recent bi-directional brain–machine interfaces being explored for neurological disorders [44–48]. As motivated earlier, the aim of this approach is to improve therapy efficacy, minimize side effects, improve battery performance, while reducing the clinical and patient burden. The prototype uses the investigational Activa™ PC + S system [49], along with the Nexus firmware update to apply embedded algorithms for stimulation therapy titration [44], both provided by Medtronic PLC. As shown in Figure 3, sensing leads were placed over the motor cortex to detect motor intentions through event-related desynchronization, while stimulation leads are placed in the thalamus to provide electrical actuation that eliminates tremors. The goal of this system was to use motor commands from the cortex to adapt thalamic stimulation in real-time, based on the subject’s immediate movement intentions. In order for the subject to exit the clinic with this adaptive system running, we applied design mitigations consistent with our proposed framework.

Starting with human factors, we establish a mental model of the control algorithm for helping assess the user model and defining potential risks. The mental model for this specific control algorithm is that the patient's movement intentions directly control the electrical stimulator's output, as illustrated in Figure 4. The intent to move is signaled by event-related changes in local field potentials that can be detected from the motor cortex; specifically, the variations of spectral power within specific frequency bands over specific regions of the brain, similar to electrocortical brain-computer interfaces [50,51]. When the perturbations of spectral oscillations signal that motor planning is in process, the control transfer element turns on the stimulation actuation; when the oscillations return to the resting state, the control transfer element turns stimulation off. As mentioned earlier, the mental model and associated risks are disease specific. For this essential tremor application, the motivation for implementing this specific control approach is to prevent tremors during intended motion, while avoiding side effects from stimulation like dysarthria, habituation and excess power consumption during rest [52].

Since the measurement of the patient's intentions will control the stimulator directly, we need to ensure a robust measuring transfer element. While the cortical signal for motion detection, the measurement element, is well characterized from classical brain-computer interface research [46–48,50,51], there is interpatient variability in specific parameters. This variability requires patient-specific tuning the band power and detection levels in the embedded comparison and control transfer elements. In addition, inpatient variability due to physiological state (e.g., awareness, sleep) was factored into the design, including the ability to fine-tune adjustments for patient-specific thresholds [49]. An additional mitigation included the use of a manual fallback mode, which included both the ability to turn the stimulator off or revert to an open-loop continuous stimulation mode that was clinically assessed to be tolerable to the patient, in other words, the fallback mode. A patient programmer allowed for this intervention at home. A final risk mitigation for the measurement transfer element was the ability to implement a continuous 'self-test' monitor to ensure the bioelectrical amplifier did not saturate. The concern is that if the amplifier saturates, then the brain oscillation that signifies the resting state would disappear and create a type I error (false positive for motion). The impact of the type I error in this case would be to hold the stimulator in a perpetual stimulation on state. As shown in Figure 5, a continuous test tone is injected into the signal chain at a discrete frequency (105 Hz) outside of the physiological bands of interest. The clinician could check that with all expected settings, this tone was preserved and that the signal chain is operating correctly. If this tone is compromised, alternative signal chain parameters are required for proper operation. Note that this approach can be implemented automatically with an embedded algorithm to generate an automated fallback mechanism out of the clinic. This automated fallback concept is similar to approaches used with cardiac pacemakers, which can sense an excessive disturbance in the signal chain, and fallback to a safe mode of open loop pacing to avoid unintended cessation of pacing therapy.

Optimizing the interaction of the adaptive stimulator and patient transfer element requires a careful balance of latency requirements. The overarching goal of the physiologic closed-loop algorithm is to prevent tremor during intended motion. A concern with a physiological control algorithm is the time it takes to turn the stimulator on when needed. If the algorithm

takes too long to respond, we create a transient type II error where breakthrough tremor might appear and be an annoyance to the patient. While if the algorithm is biased to be highly sensitive to respond quickly, the benefits of adaption might be lost due to type I errors. Latency in this therefore implicitly reflected in the trade-off between sensitivity and specificity. Since today's open-loop systems are 100% sensitive and relatively nonspecific, essentially on all the time, the clinician can err on the side of providing acceptable latency at the expense of the occasional false positive for motion intention (e.g., bias for type I errors vs type II errors). The selection of the measurement approach also presents performance and risk trade-offs. For example, a potential advantage of cortical sensing over wearable accelerometers is that the latency of the sensing is reduced as cortical intention signals precede overt motion [51], providing a head start to stimulation ramping. However, cortical sensing has the trade-off of requiring another electrode in the system [45].

Latency also factors into the design of the control transfer element's actuation limits for both the amplitude and the transition rate for stimulation. Specifically, the control of thalamic stimulation should be bounded to set levels to prevent adverse side effects. Of concern for thalamic stimulation is an overly rapid stimulation-ramping rate or excessive stimulation, either of which might drive side effects [52,53]. To prevent these occurrences, the clinician programming computer defines the safe operating region in the clinic, under direct observation. These parameters can then be downloaded to the device as protected variables that set hard limits for the embedded algorithm. If for some reason the embedded algorithm attempts to exceed these limits, the control element's explicit command value is ignored, and the predefined, clinician-specified limit value provided to the stimulator output is used instead. Similar algorithm limitations restrict the ramping rate of stimulation to avoid undesirable paresthesia.

Even with these control transfer element limits in place, we still implemented a fallback mode as a final risk mitigation in case of an unforeseen issue. With a finite probability that the patient's biomarker variability or situation might lead to inappropriate algorithm performance, the bidirectional BMI uses a patient controller with a preloaded stimulation program defined by the clinician as 'safe' open-loop stimulation. Before the embedded algorithm can operate independently, the implant verifies that this program is configured. Once operational, a patient can press a button that immediately diverts to this fallback state. In addition, an interrogation by any clinician programmer also drives the embedded algorithm to this fallback mode. This mitigation ensures that in an emergent situation a clinician naive to the research is dealing with standard neurostimulator operation that aligns with their own mental model of how the typical system should operate. Once the patient is ready and the hazard no longer present, they can use their programmer to reenter the automated feedback state. In the future, there might be additional reasons to enhance the fallback feature. For example, systems that integrate wearable sensors or apply distributed algorithms might be sensitive to dropped telemetry packets. When this occurs, the algorithm might be unable to assess the appropriate action to take; for these cases, the system can enter the fallback state until the connection is reestablished, and sensing is recovered.

The patient's handheld programmer is also useful for providing state indicators and can provide a user interface for patients to interact with the system. The programmer conveys the

state of the control element, clearly differentiating whether the system is in fallback mode or running in adaptive stimulation mode. However, state indicators can also employ direct physiological feedback through imbedded sensory feedback. For example, in pilot experiments, some patients preferred having a small amount of transient paresthesia – the side-effect sensation from rapid stimulation ramp rates indicating stim turning on and off – as means to signal a transition state is occurring in the device. The patients can use this sensation as an indicator that the algorithm is in transition, which provides assurance that the algorithm is ramping up stimulation as they initiate a task. This type of sensory feedback through stimulation could prove useful for many brain–machine interface applications, and is an active area of inquiry [54,55].

Finally, to support training and fault analysis, we used embedded data logs configured in the implant. Configuring the algorithm requires mapping the signals from the cortex to patient state. Telemetry streaming allows for uploading signals in the four states of stim on/off combined with motion on/off, to allow for calibration of the algorithm for embedded operation. In addition, patient-triggered data recorders allow for gathering snapshots of data when the algorithm is not functioning correctly or patients are showing other symptoms [46]. These embedded data logs provide the ability to refine the algorithm based on faults or disturbances in the home environment [47]. Finally, the prototyping environment allows for the clinician to train and validate the algorithm in the clinic using a computer-in-the-loop, providing confidence that the algorithm will function properly when the physiologic algorithm is embedded in the device before allowing the patient to leave the clinic for at-home testing [3].

The methods described in this section allowed for execution of a clinical study with an automated, adaptive brain stimulator with subjects out of the clinic for 6 months. The results of this study are currently being compiled, but the relevant observation for this perspective is that no adverse events were reported during the protocol. In addition, the methods described here are generalizable to other disease states, and were applied for Tourette disorder [56], Parkinson’s disease [47], and are currently being applied to epilepsy systems [57].

Summary

This technical perspective provided an overview of design methods for intelligent bioelectronic systems based on a physiologic closed-loop model of control. While implementing these design methods does not guarantee safety and reliability, by following a systematic process of design control for each specific use case, identifying risk and design mitigations the designer can maximize the robustness of the intelligent system in real-world deployment. The design principles developed in this brief are quite generalizable, as demonstrated by the historical examples and reinforced by our case study exploring the design of a bidirectional brain–machine interface operating as an automated, adaptive deep brain stimulator; similar ideas guide the design for state-of-the-art automated anesthesia machine (anesthesia) and ‘artificial pancreas’ devices [5,14,30].

Conclusion

In conclusion, the system designer needs to plan intentionally for limiting the potential downside of intelligent medical systems, while also exploiting the upsides, especially when exploring novel technology or novel scientific concepts. To paraphrase the English statesman Benjamin Disraeli, “Be prepared for the worst, but hope for the best.”

Future perspective

The field of ‘smart’ biomedical systems is fuelled by the need for improved therapies delivered with less clinical and economic burden. Complementary market forces such as consumer electronics, large-scale algorithm processing and data sciences, increased investments in medical research, and social demographics all catalyse the shift to more capable medical devices. We foresee the continued integration of implants and wearables that rely on both local and distributed computing resources. Such systems will allow for digital phenotyping replacing or augmenting in-clinic practice, and provide an objective method for tracking disease. Platform systems with these capabilities will drive ‘device-enabled precision medicine’ through physiology-based biomarker discovery for neurological and psychiatric diseases gathered in natural settings. Precision medicine methods applied to devices should enable patient-specific adaptive stimulation, and accelerate therapy optimization while reducing care burden.

We also see the need for improved design mitigations in the future, as systems become more automated. First, the need for enhanced training: as these systems become ubiquitous, more systematic clinician training becomes crucial. Recent examples from transportation highlight the need for intuitive mental models that the user can rely on when using automated controls, and that guide intervention if a patient is at risk. Second, engaging patient feedback: as systems are deployed into natural settings, there is a need for scalable systems where patients and consumers can provide feedback to clinicians on control system performance. This feedback will enhance the training set for improving performance from the user’s perspective. Such data can enable a third area for improvement, algorithm design: as these systems expand into multi-modes of actuation and sensing, clinicians may be able to optimize more than one clinical parameter (e.g., in Parkinson’s disease, alleviating both hyperkinetic and hypokinetic movement while protecting cognition). This level of design motivates a control layer in ‘parameter’ space, and considerations of additional configuration aids using techniques from reinforcement learning that help automate the parameter selection process in light of clinician and patient goals. Finally, the expansion of therapies into mood disorders and anxiety prompt the need for consideration of the ethics of automated neural control, with special attention to informed consent and transparency in protocol development.

References

Papers of special note have been highlighted as: • of interest; •• of considerable interest

1. Denison T, Morris M, Sun F. Building a bionic nervous system. *IEEE Spect.* 2015; 52(2):32–39. [**Provides an overview of the state-of-the-art in closed-loop medical devices for the nervous system.**]
2. Famm K, Litt B, Tracey KJ, Boyden ES, Slaoui M. Drug discovery: a jump-start for electroceuticals. *Nature.* 2013; 496(7444):159–161. [PubMed: 23579662] [**Provides an vision for future uses of bioelectronic medicines in the periphery.**]
3. Afshar P, Khambhati A, Stanslaski S, et al. A translational platform for prototyping closed-loop neuromodulation systems. *Front Neural Circuits.* 2012; 6:117. [PubMed: 23346048]
4. Wagner FB, Mignardot J-B, Le Goff-Mignardot CG, et al. Targeted neurotechnology restores walking in humans with spinal cord injury. *Nature.* 2018; 563(7729):65–71. [PubMed: 30382197]
5. Majeed W, Thabit H. Closed-loop insulin delivery: current status of diabetes technologies and future prospects. *Expert Rev Med Devices.* 2018; 15(8):579–590. [PubMed: 30027775]
6. Parasuraman R, Sheridan TB, Wickens CD. A model for types and levels of human interaction with automation. *IEEE Trans Syst Man Cybern A Syst Hum.* 2000; 30(3):286–297. [PubMed: 11760769]
7. Benabid AL, Pollak P, Gervason C, et al. Long-term suppression of tremor by chronic stimulation of the ventral intermediate thalamic nucleus. *Lancet.* 1991; 337(8738):403–406. [PubMed: 1671433]
8. Volkmann J, Herzog J, Kopper F, Deuschl G. Introduction to the programming of deep brain stimulators. *Mov Disord.* 2002; 17(Suppl. 3):S181–S187. [PubMed: 11948775]
9. Moro E, Allert N, Eleopra R, et al. A decision tool to support appropriate referral for deep brain stimulation in Parkinson's disease. *J Neurol.* 2009; 256(1):83–88. [PubMed: 19221846]
10. International Organization for Standardization. Medical electrical equipment - part 1-10: general requirements for basic safety and essential performance - collateral standard: requirements for the development of physiologic closed-loop controllers. IEC 60601-1-10:2007 www.iso.org/obp/ui/#iso:std:iec:60601:-1-10:ed-1:v1:en,fr
11. Widge AS, Malone DA, Dougherty DD. Closing the loop on deep brain stimulation for treatment-resistant depression. *Front Neurosci.* 2018; 12:175. [PubMed: 29618967]
12. Den Dulk K, Bouwels L, Lindemans F, Rankin I, Brugada P, Wellens HJ. The activitrax rate responsive pacemaker system. *Am J Cardiol.* 1988; 61(1):107–112. [PubMed: 3336998]
13. Lloyd M, Reynolds D, Sheldon T, et al. Rate adaptive pacing in an intracardiac pacemaker. *Heart Rhythm.* 2017; 14(2):200–205. [PubMed: 27871854] [**Provides an example of automaticity from the cardiac pacing space.**]
14. Dadlani V, Pinsker JE, Dassau E, Kudva YC. Advances in closed-loop insulin delivery systems in patients with type 1 diabetes. *Curr Diab Rep.* 2018; 18(10):88. [PubMed: 30159816] [**Provides an overview of recent advancements in the 'artificial pancreas'.**]
15. Zavitsanou S, Mantalaris A, Georgiadis MC, Pistikopoulos EN. *In silico* closed-loop control validation studies for optimal insulin delivery in type 1 diabetes. *IEEE Trans Biomed Eng.* 2015; 62(10):2369–2378. [PubMed: 25935026]
16. Appelboom G, Camacho E, Abraham ME, et al. Smart wearable body sensors for patient self-assessment and monitoring. *Arch Public Health.* 2014; 72(1):28–28. [PubMed: 25232478]
17. Leelarathna L, English SW, Thabit H, et al. Accuracy of subcutaneous continuous glucose monitoring in critically ill adults: improved sensor performance with enhanced calibrations. *Diabetes Technol Ther.* 2014; 16(2):97–101. [PubMed: 24180327]
18. Cote GL, Lec RM, Pishko MV. Emerging biomedical sensing technologies and their applications. *IEEE Sens J.* 2003; 3(3):251–266.
19. Puers R, De Bruyker D, Cozma A. A novel combined redundant pressure sensor with self-test function. *Sens Actuators A Phys.* 1997; 60(1):68–71.
20. Schwiebert, L; Gupta, SKS; Weinmann, J. Research challenges in wireless networks of biomedical sensors. *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*; Rome, Italy: 2001 Jul 16-21.
21. Boughton CK, Homeworck R. Is an artificial pancreas (closed-loop system) for type 1 diabetes effective? *Diabet Med.* 2018; 36(3):279–286. [PubMed: 30183096]
22. Elleri D, Maltoni G, Allen JM, et al. Safety of closed-loop therapy during reduction or omission of meal boluses in adolescents with type 1 diabetes: a randomized clinical trial. *Diabetes Obes Metab.* 2014; 16(11):1174–1178. [PubMed: 24909206]

23. Gillberg J. Detection of cardiac tachyarrhythmias in implantable devices. *J Electrocardiol.* 2007; 40(6 Suppl):S123–S128. [PubMed: 17993309]
24. Shome S, Li D, Thompson J, McCabe A. Improving contemporary algorithms for implantable cardioverter-defibrillator function. *J Electrocardiol.* 2010; 43(6):503–508. [PubMed: 20719333]
25. Geller EB, Skarpaas TL, Gross RE, et al. Brain-responsive neurostimulation in patients with medically intractable mesial temporal lobe epilepsy. *Epilepsia.* 2017; 58(6):994–1004. [PubMed: 28398014]
26. Sun FT, Morrell MJ. The RNS System: responsive cortical stimulation for the treatment of refractory partial epilepsy. *Expert Rev Med Devices.* 2014; 11(6):563–572. [PubMed: 25141960]
27. Sun FT, Morrell MJ. Closed-loop neurostimulation: the clinical experience. *Neurotherapeutics.* 2014; 11(3):553–563. [PubMed: 24850309] [**Provides an overview of the clinical point-of-view for emerging smart neural devices.**]
28. Baldassano S, Wulsin D, Ung H, et al. A novel seizure detection algorithm informed by hidden Markov model event states. *J Neural Eng.* 2016; 13(3):036011–036011. [PubMed: 27098152]
29. Garg SK, Weinzimer SA, Tamborlane WV, et al. Glucose outcomes with the in-home use of a hybrid closed-loop insulin delivery system in adolescents and adults with type 1 diabetes. *Diabetes Technol Ther.* 2017; 19(3):155–163. [PubMed: 28134564]
30. Grosman B, Ilany J, Roy A, et al. Hybrid closed-loop insulin delivery in type 1 diabetes during supervised outpatient conditions. *J Diabetes Sci Technol.* 2016; 10(3):708–713. [PubMed: 26880389]
31. Dawson TW, Caputa K, Stuchly MA, Kavet R. Pacemaker interference by 60-Hz contact currents. *IEEE Trans Biomed Eng.* 2002; 49(8):878–886. [PubMed: 12148827]
32. Dodinot B, Godenir JP, Costa AB. Electronic article surveillance: a possible danger for pacemaker patients. *Pacing Clin Electrophysiol.* 1993; 16(1 Pt 1):46–53. [PubMed: 7681175]
33. Gordon RS, O'Dell KB. Permanent internal pacemaker safety in air medical transport. *J Air Med Transp.* 1991; 10(2):22–23.
34. Marco D, Eisinger G, Hayes DL. Testing of work environments for electromagnetic interference. *Pacing Clin Electrophysiol.* 1992; 15(11 Pt 2):2016–2022. [PubMed: 1279591]
35. Shah AD, Morris MA, Hirsh DS, et al. Magnetic resonance imaging safety in nonconditional pacemaker and defibrillator recipients: a meta-analysis and systematic review. *Heart Rhythm.* 2018; 15(7):1001–1008. [PubMed: 29458192]
36. Irnich W, Steen-Mueller MK. Pacemaker sensitivity to 50 Hz noise voltages. *Europace.* 2011; 13(9):1319–1326. [PubMed: 21551474]
37. Becker R, Ruf-Richter J, Senges-Becker JC, et al. Patient alert in implantable cardioverter defibrillators: toy or tool? *J Am Coll Cardiol.* 2004; 44(1):95–98. [PubMed: 15234415]
38. Purerfellner H, Sanders P, Pokushalov E, et al. Miniaturized reveal LINQ insertable cardiac monitoring system: first-in-human experience. *Heart Rhythm.* 2015; 12(6):1113–1119. [PubMed: 25728756]
39. Reinsch N, Ruprecht U, Buchholz J, Diehl RR, Kalsch H, Neven K. The BioMonitor 2 insertable cardiac monitor: clinical experience with a novel implantable cardiac monitor. *J Electrocardiol.* 2018; 51(5):751–755. [PubMed: 30177307] [**Highlights the use of telemedicine and implantable sensors for remote patient management.**]
40. Jung W, Rillig A, Birkemeyer R, Miljak T, Meyerfeldt U. Advances in remote monitoring of implantable pacemakers, cardioverter defibrillators and cardiac resynchronization therapy systems. *J Interv Card Electrophysiol.* 2008; 23(1):73–85. [PubMed: 18821006]
41. Parasuraman R, Molloy R, Singh IL. Performance consequences of automation-induced 'complacency'. *Int J Aviation Psychol.* 1993; 3(1):1–23. [**Considers the unintended side effects of automaticity in systems.**]
42. Cook MJ, O'Brien TJ, Berkovic SF, et al. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. *Lancet Neurol.* 2013; 12(6):563–571. [PubMed: 23642342]
43. Herron JA, Thompson MC, Brown T, Chizeck HJ, Ojemann JG, Ko AL. Cortical brain-computer interface for closed-loop deep brain stimulation. *IEEE Trans Neural Syst Rehabil Eng.* 2017; 25(11):2180–2187. [PubMed: 28541211]

44. Khanna P, Swann NC, De Hemptinne C, et al. Neurofeedback control in Parkinsonian patients using electrocorticography signals accessed wirelessly with a chronic, fully implanted device. *IEEE Trans Neural Syst Rehabil Eng.* 2017; 25(10):1715–1724. [PubMed: 28113590]
45. Swann NC, De Hemptinne C, Miocinovic S, et al. Chronic multisite brain recordings from a totally implantable bidirectional neural interface: experience in 5 patients with Parkinson's disease. *J Neurosurg.* 2018; 128(2):605–616. [PubMed: 28409730]
46. Swann NC, De Hemptinne C, Miocinovic S, et al. Gamma oscillations in the hyperkinetic state detected with chronic human brain recordings in Parkinson's Disease. *J Neurosci.* 2016; 36(24):6445–6458. [PubMed: 27307233]
47. Swann NC, De Hemptinne C, Thompson MC, et al. Adaptive deep brain stimulation for Parkinson's disease using motor cortex sensing. *J Neural Eng.* 2018; 15(4)
48. Leuthardt EC, Schalk G, Roland J, Rouse A, Moran DW. Evolution of brain-computer interfaces: going beyond classic motor physiology. *Neurosurg Focus.* 2009; 27(1):e4.
49. Rouse AG, Stanslaski SR, Cong P, et al. A chronic generalized bi-directional brain-machine interface. *J Neural Eng.* 2011; 8(3)
50. Leuthardt EC, Schalk G, Wolpaw JR, Ojemann JG, Moran DW. A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng.* 2004; 1(2):63–71. [PubMed: 15876624]
51. Wolpaw JR, Birbaumer N, McFarland DJ, Pfurtscheller G, Vaughan TM. Brain-computer interfaces for communication and control. *Clin Neurophysiol.* 2002; 113(6):767–791. [PubMed: 12048038]
52. Gibson WS, Jo HJ, Testini P, et al. Functional correlates of the therapeutic and adverse effects evoked by thalamic stimulation for essential tremor. *Brain.* 2016; 139(Pt 8):2198–2210. [PubMed: 27329768]
53. Swan BD, Gasperson LB, Krucoff MO, Grill WM, Turner DA. Sensory percepts induced by microwire array and DBS microstimulation in human sensory thalamus. *Brain Stimul.* 2018; 11(2):416–422. [PubMed: 29126946]
54. Bowsher K, Civillico EF, Coburn J, et al. Brain-computer interface devices for patients with paralysis and amputation: a meeting report. *J Neural Eng.* 2016; 13(2)
55. Graczyk EL, Resnik L, Schiefer MA, Schmitt MS, Tyler DJ. Home use of a neural-connected sensory prosthesis provides the functional and psychosocial experience of having a hand again. *Sci Rep.* 2018; 8(1):9866. [PubMed: 29959334]
56. Shute JB, Okun MS, Opri E, et al. Thalamocortical network activity enables chronic tic detection in humans with Tourette syndrome. *NeuroImage Clin.* 2016; 12:165–172. [PubMed: 27419067]
57. Kremen V, Brinkmann BH, Kim I, et al. Integrating brain implants with local and distributed computing devices: a next generation epilepsy management system. *IEEE J: Trans Eng Health Med.* 2018; 6:1–12. [**Provides an overview of a cloud connected, distributed algorithm for epilepsy management, which might become more common with future medical devices.**]

Executive summary

The motivation for 'smart' systems

- Medical devices are beginning to use sensors and algorithms to automate the delivery of therapies.
- Automaticity has the potential to improve therapy outcomes, reduce side effects and offload clinical and patient burden.
- However, automaticity also brings new risks that need to be addressed by the system designer.

Frameworks to guide 'smart' system design

- Regulations exist for physiologic controllers, and apply to many automated closed-loop systems such as incubators and syringe pumps.
- The guidance documents for these systems are also useful for guiding implantable medical devices and systems.
- The application of a physiologic control framework provides a common design language for designing automated bioelectronic systems and managing risks.

Building intuition: mapping historical examples of risk mitigation to the framework

- Existing medical devices with embedded automaticity provide specific examples of how the framework can be applied in practice.
- The cardiac pacemaker, artificial pancreas and neuromodulators provide implicit examples of design choices that align with the intent of regulatory guidance, which can be made explicit through the use of the design framework.
- These predicate designs provide examples of both core design elements and key risk considerations and mitigations for single-fault failures.

An integrated 'smart' system case study: bidirectional brain-computer interfacing for treating essential tremor

- An exemplar design from a state-of-the-art investigational study is presented as an example of a system that was designed within the framework.
- Specific design steps and risk mitigations are described that highlight key elements of the physiologic control framework.
- The brain-interfacing example integrates all elements of the framework into a single unified system, which can be used as a reference template for bioelectronic medical systems with automaticity and feedback.

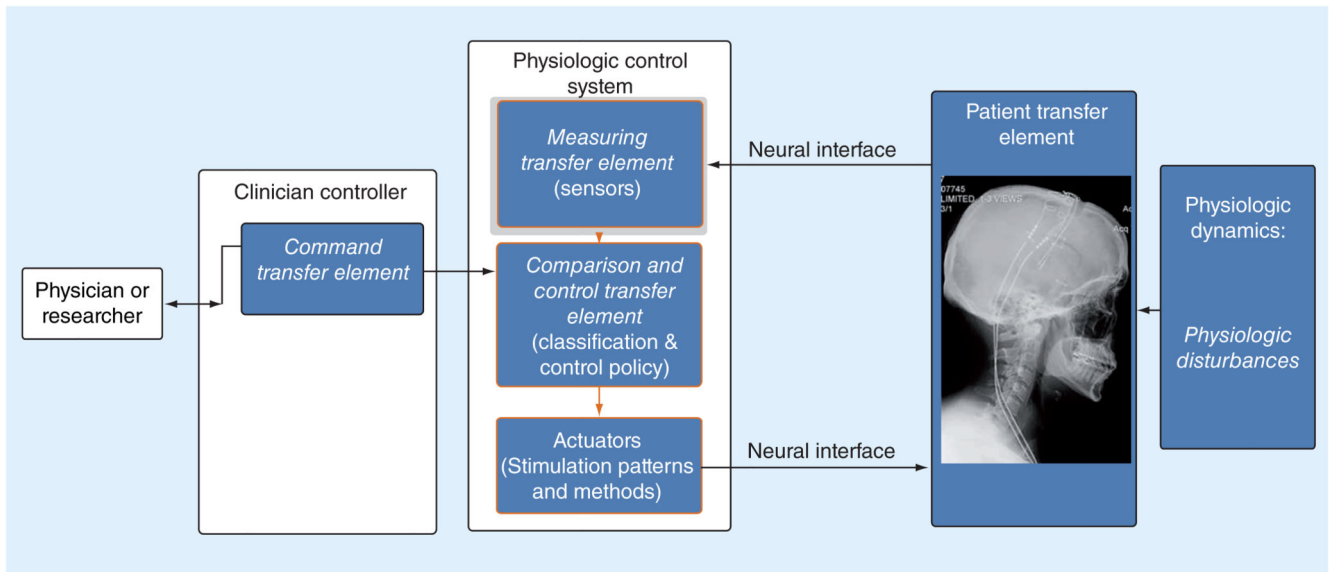


Figure 1. The base model for an ‘intelligent’ bioelectronic system using the nomenclature of the International Electrotechnical Commission standard 60601-1-10 for establishing safety and essential performance in physiologic closed-loop systems.

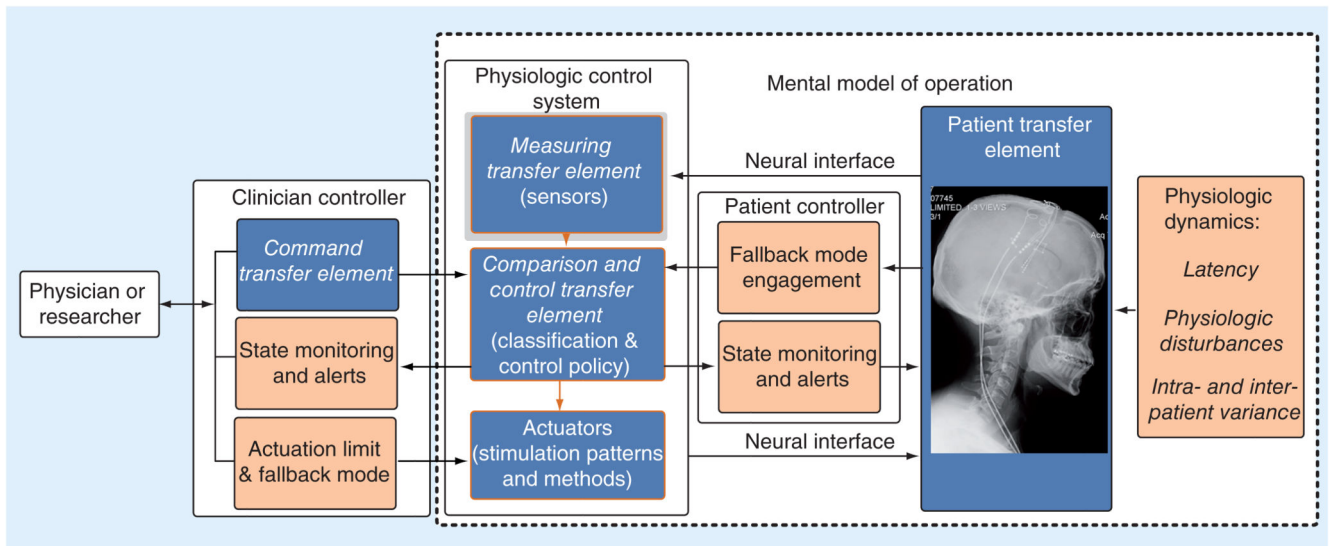


Figure 2. An enhanced model for an 'intelligent' bioelectronic system using the nomenclature of the International Electrotechnical Commission standard 60601-1-10 for establishing safety and essential performance in physiologic closed-loop systems; risk mitigations are added with the brown boxes.

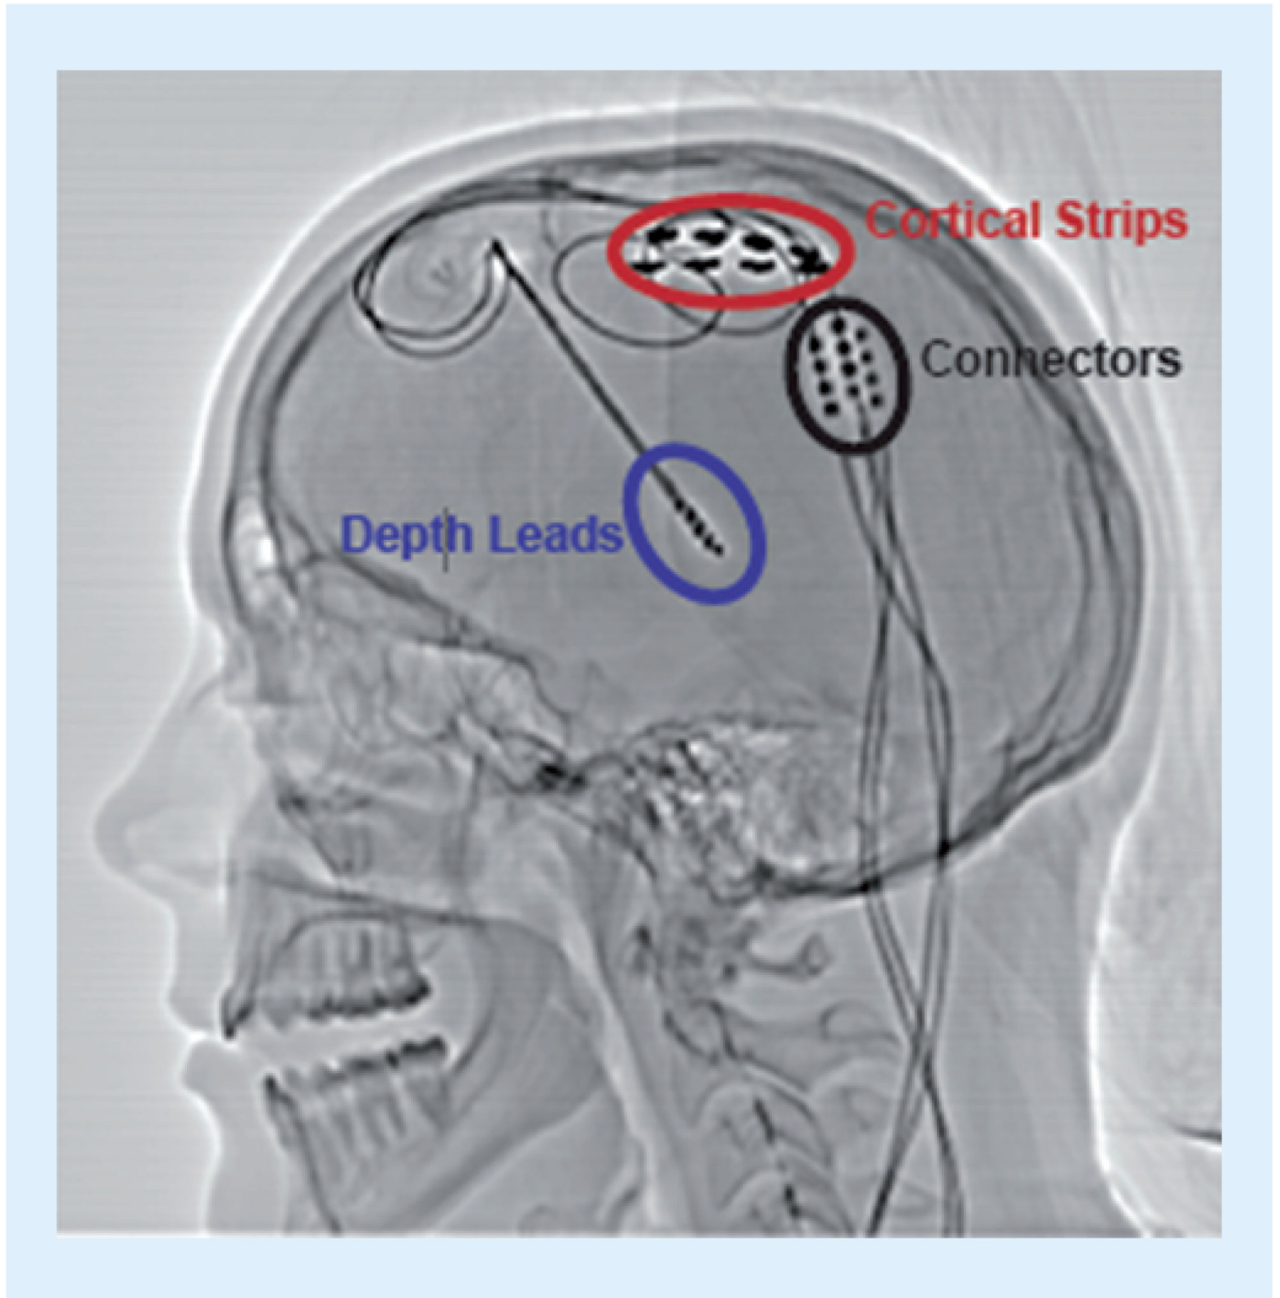


Figure 3. Cortical–thalamic lead placement for an adaptive neurostimulator. Cortical sensing electrodes (measurement element) are circled in red, while depth stimulation electrodes in the thalamus (actuation) are circled in blue.




Figure 4. The mental model for the physiologic controller using cortical sensing to drive a thalamic lead – when the patient ‘thinks’ about motion to pour a cup of coffee, the stimulator turns on; when he stops thinking about motion, the stimulator turns off.
By limiting stimulation to the motion state, therapy side effects and battery usage might be better managed.

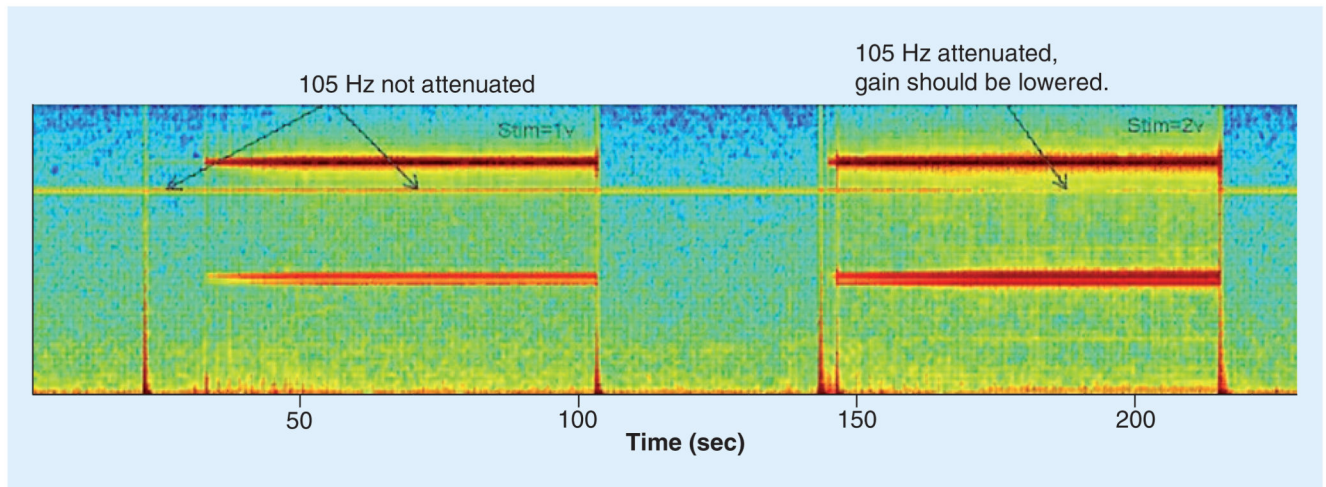


Figure 5. The use of a test tone to monitor the integrity of the measurement transfer element (biopotential amplifier).

The 105 Hz tone is injected in parallel to the physiological signal and should always be present. If it disappears due to a combination of excessive channel mismatch, stimulation gain and amplification, the tone is compromised and corrective actions should be taken, such as turning down the gain or entering a fallback mode. This example was provided by Medtronic PLC for training purposes.