

Rapid communication

ADAPTIVE CONTROL OF EPILEPTIFORM EXCITABILITY IN AN *IN VITRO* MODEL OF LIMBIC SEIZURES.

Gabriella Panuccio^{a*}, Arthur Guez^{b*}, Robert Vincent^c, Massimo Avoli^a and Joelle Pineau^{c,**}

^aMontreal Neurological Institute and Department of Neurology & Neurosurgery, McGill University, 3801 University Street, Montreal, QC, H3A 2B4, Canada; ^bGatsby Unit, UCL, London, UK, WC1N 3AR; ^cSchool of Computer Science, McGill University, 3480 University Street, Montreal, QC, H3A 2A7, Canada.

Keywords: Closed-loop strategy, Deep brain stimulation, Epilepsy

* These authors contributed equally

****Correspondence to:** J. Pineau, Ph.D.
3480 University Street, Montreal, QC
H3A 2A7, Canada
Tel (514) 398 7071
Fax (514) 398 3883
e-mail: jpineau@cs.mcgill.ca

Abstract

Deep Brain Stimulation (DBS) is a promising tool for treating drug-resistant epileptic patients. Currently, the most common approach is fixed-frequency stimulation (periodic pacing) by means of stimulating devices that operate under open-loop control. However, a drawback of this DBS strategy is the impossibility of tailoring a personalized treatment, which also limits the optimization of the stimulating apparatus. Here, we propose a novel DBS methodology based on a closed-loop control strategy, developed by exploiting statistical machine learning techniques, in which stimulation parameters are adapted to the current neural activity thus allowing for seizure suppression that is fine-tuned on the individual scale (adaptive stimulation). By means of field potential recording from adult rat hippocampus-entorhinal cortex (EC) slices treated with the convulsant 4-aminopyridine we determined the effectiveness of this approach compared to low-frequency periodic pacing, and found that the closed-loop stimulation strategy: (i) has similar efficacy as low-frequency periodic pacing in suppressing ictal-like events but (ii) is more efficient than periodic pacing in that it requires less electrical pulses. We also provide evidence that the closed-loop stimulation strategy can alternatively be employed to tune the frequency of a periodic pacing strategy. Our findings indicate that the adaptive stimulation strategy may represent a novel, promising approach to DBS for individually-tailored epilepsy treatment.

¹ Abbreviations

1 DBS: Deep Brain Stimulation; EC: Entorhinal Cortex.

Introduction

Epilepsy is a highly-prevalent chronic neurological disorder (Kotsopoulos, et al., 2002) and up to 30% of epilepsy patients do not respond to pharmacological treatment (Kwan and Brodie, 2000). Moreover, resective surgery comes at the risk of physical and cognitive impairments. An emerging alternative treatment for drug-resistant patients is deep brain stimulation (DBS), i.e. direct delivery of electrical pulses to the brain, aiming at modulating neuronal excitability and thus halting or even preventing seizures. Recent studies have been evaluating the effectiveness of DBS strategies in drug-resistant epileptic patients (Boon, et al., 2007; Ellis and Stevens, 2008; Hamani, et al., 2009; Marks, 2008). However, conclusive evidence in humans is difficult to achieve, due to the large variance in the disease and symptoms among patients; it is therefore difficult to choose a single stimulation protocol that is effective on a collection of individuals.

There are many parameters to select when applying DBS, including the target area for stimulation, as well as the frequency, intensity and pattern of the electrical pulses. These parameters can be specified either through an open-loop paradigm, or as a closed-loop control system. An open-loop strategy uses preset parameters to deliver stimulation, without monitoring electrical cortical activity. Examples of open-loop strategies include fixed-frequency stimulation (periodic pacing), as well as stimulation strategies based on a fixed random process (e.g., gaussian noise generator) (Durand and Bikson, 2001). In a closed-loop strategy, the stimulation parameters are dynamically changed in response to sensor readings of brain activity. This can be achieved by using software to automatically detect an impending seizure and administering a fixed stimulation protocol designed to terminate the seizure (Cohen-Gadol, et al., 2003; Durand and Bikson, 2001; Fountas, et al., 2005; Kossoff, et al., 2004; Loscher and Schmidt, 2004; Osorio, et al., 2005; Theodore and Fisher, 2004). This can also be achieved through more sophisticated feedback control methods (Guez, et al., 2008; Pineau, et al., 2009).

More substantial evidence for the ability of DBS to successfully reduce epileptic symptoms has been produced using animal models of epilepsy. A number of studies focused on finding the appropriate stimulation site and frequency of stimulation for open-loop (D'Arcangelo, et al., 2005; Durand and Bikson, 2001; Ellis and Stevens, 2008; Schiller and Bankirer, 2007) as well as for closed-loop strategies (Bush and Pineau, 2009; Durand and Bikson, 2001; Nakagawa and Durand, 1991; Schiff, et al., 1994; Schiller and Bankirer, 2007). By using an *in vitro* model of limbic ictogenesis, we have previously reported that repetitive low-frequency stimulation, delivered in the subiculum at frequencies similar to those of the CA3-driven interictal discharges, decreases epileptiform synchronization in the entorhinal cortex (EC) (Barbarosie and Avoli, 1997). In particular, we have shown that the 1 Hz frequency exhibits maximal efficacy in reducing EC ictogenesis (D'Arcangelo, et al., 2005).

We have put forward, in previous studies, the use of statistical learning techniques to automatically optimize closed-loop strategies. We show how the strategies can be learned from field potential recordings in rat brain slices in which epileptiform discharges were induced by superfusion with the convulsant 4-aminopyridine (4AP) (Guez, et al., 2008; Pineau, et al., 2009). These studies suggested that closed-loop DBS could achieve successful suppression of ictal-like activity using less pulses than open-loop periodic pacing strategies; however these results were obtained using an *in silico* model of epilepsy. Here, we have evaluated such statistical learning closed-loop strategies *in vitro* by using the 4AP model of limbic ictogenesis in adult rat hippocampus-EC slices. Our findings show that the learned controller is able to perform as well as the 1 Hz open-loop paradigm, while reducing the total stimulation delivered in most slices. In contrast, applying periodic pacing at the same effective mean frequency found by the closed-loop controller is not as effective in suppressing ictal-like activity as our adaptive algorithm's solution.

Methods

Brain slice preparation and maintenance

All efforts were made to minimize the number of animals used and their suffering. All the procedures were carried on in accordance to the Canadian Council on Animal Care and McGill University guidelines. Nineteen male, adult Sprague-Dawley rats (250-300 g) were decapitated under deep isoflurane anesthesia. The brain was quickly removed and placed in cold (0-2° C) artificial cerebro-spinal fluid (ACSF), continuously bubbled with gas mixture (CO₂ 5% and O₂ 95%) to equilibrate at pH=7.35-7.40, and having the following composition (mM): 124 NaCl, 2 KCl, 2 MgSO₄, 2 CaCl₂, 1.25 KH₂PO₄, 26 NaHCO₃ and 10 D-glucose. *Partially disconnected* combined hippocampus-EC slices (450 μm thick) including the most ventral part of the hippocampal formation were cut as previously described (Panuccio, et al., 2010) using a VT1000S vibratome (Leica, Germany). In these brain slices fast CA3-driven interictal-like activity disclosed by 4AP application was observed within the hippocampus proper only, i.e. it did not propagate to the EC (*cf.*, Avoli, et al., 1996, but see also Avoli, et al., 2002). Slices were then transferred to an interface recording chamber, lying between warm (~ 32° C) ACSF and humidified gas (CO₂ 5% and O₂ 95%), where they were allowed to recover for at least 1 hour before beginning continuous bath-application of 4AP. Slices were continuously perfused at ~1 ml/min.

Field potential recording and stimulation paradigms

Field potential recordings were made with ACSF-filled pipettes (tip diameter <10 μm; resistance = 5-10 MΩ) pulled from borosilicate capillary tubing (World Precision Instruments Inc., Sarasota, FL, USA) using a P-97 puller (Sutter Instrument, Novato, CA, USA). Extracellular signals were fed to a Cyberamp 380 amplifier (Molecular Devices, Palo Alto, CA) connected to a digital interface device (Digidata 1320A, Molecular Devices). Data were acquired at a sampling rate of 5 KHz (low-pass filtered at 2 KHz), using the software Clampex 8.2 (Molecular Devices), stored on the hard drive and analyzed off-line using pClamp 9.0. Recording electrodes were placed in the deep layers of the medial EC, CA3 stratum pyramidale or radiatum, and the pyramidal layer of the proximal subiculum (see Fig. 1A). Extracellular current pulses (0.1-2.25 mA, pulse width 100 μs) were delivered in the pyramidal layer of the proximal subiculum through a bipolar concentric Pt-Ir electrode (FHC, Bowdoin, ME, USA) plugged onto a high voltage stimulus isolator unit (A360, WPI Inc., Sarasota, Florida, USA) connected to the pulse generator Pulsemaster A300 (WPI Inc., Sarasota, Florida, USA). Current intensity required to induce a response with failure rate <80% was established before the beginning of any stimulation protocol and was kept constant throughout the experiment. The Cyberamp 380 amplifier was also connected to another digital interface device (USB-6221 M Series, National Instruments, Texas, USA) that acquired data at a sampling rate of 5 KHz using an in-house software. That software was performing signal processing in real time and was querying the adaptive controller software for the stimulation actions. Those stimulation requests were then transmitted digitally from the National Instruments digitizer to the pulse generator.

Before running the adaptive code, we assessed slice viability in terms of subicular projections to the EC by testing the efficacy of periodic pacing at 1 Hz, which has proved to be the most effective low-frequency periodic pacing protocol in terms of suppression of ictal activity (*cf.* D'Arcangelo et al., 2005). During the training phase, three different stimulation strategies were applied in 4 slices as follows: (i) periodic pacing at 0.5Hz, (ii) period pacing at 1Hz, (iii) periodic pacing at 2 Hz. An early version of the closed-loop controller was also applied to other 4 slices to obtain more relevant training data. Each stimulation paradigm was preceded by a control period and followed by a recovery phase, which also served as the control of the following stimulation protocol. In total, about 12 hours of recording (including control and stimulation

protocols) were acquired to train our closed-loop controller. During the validation phase, three different stimulation strategies were applied on 11 slices as follows: (i) periodic pacing at 1 Hz, (ii) adaptive stimulation using the closed-loop controller, and (iii) periodic pacing at the average frequency of the adaptive stimulation protocol (we call this the *effective frequency stimulation*). As above, each stimulation paradigm was preceded by a control period and followed by a recovery phase, which also served as the control of the following stimulation protocol. The effective frequency stimulation, f (Hz), is a periodic pacing strategy. A different effective frequency is computed for each slice after application of the adaptive controller using the following equation:

$$f = \frac{n_s}{T}$$

where n_s is the total number of pulses delivered during the duration T , in seconds, of the adaptive stimulation protocol. This strategy effectively includes as many pulses as the adaptive strategy applied to the slice, but distributed in a periodic manner.

Adaptive stimulation design

The stimulation patterns for the adaptive stimulation strategy are computed in real-time based on a function relating the observed neural activity and optimal stimulation parameters. This function is learned *a priori* using data collected from the training slices. This function defines the parameters of the adaptive stimulation strategy over the full range of observed neural conditions. The function is then used on the validation slices, to match observed neural signal with optimal stimulation parameters. A full description of the mathematical and computational methods underlying the adaptive controller can be found in Guez, et al. (2008) and Pineau, et al. (2009). Here, we briefly summarize the mathematical method.

Field potential recordings were processed using Fast Fourier Transforms over different window lengths to extract spectral features forming the state vector s on which the adaptive controller based its stimulation decision at each time step. A cost c was associated with performing a stimulation action a in a state s , whose cost was mainly influenced by the presence of epileptiform activity in s but also by the frequency of stimulation described by a .

The controller could then choose, by the mean of an action a at a state s , between not stimulating or stimulating at either 0.5 Hz, 1 Hz, or 2 Hz for the duration of the next reference window. The goal of the controller was to reduce the long-term accumulation of those costs. To achieve that, a sophisticated regression tool, called extremely randomized trees Geurts, et al., 2006, was used to learn the long-term cumulative cost $Q(s,a)$ of applying an action a in any state s using the field potential recordings in the training dataset. In the validation phase, an optimal stimulation action a_t , can then be selected based on the current neuronal network state s_t , using the best learned controller function $Q(s_t, a_t) = \max_a Q(s_t, a)$.

Data and statistical analysis

We arbitrarily defined ictal-like (hereafter termed ictal) discharges to be those epileptiform events resembling EEG ictal activity and lasting longer than 3 s. We used the following parameters as performance indicators of the adaptive controller:

(i) $t_p^i \in [0,1]$, which denotes the proportion of ictal time during protocol p for the slice i , where for each slice the following protocols are executed in the following sequence: Control (CTRL1) \rightarrow 1 Hz stimulation \rightarrow Recovery (CTRL2) \rightarrow Adaptive stimulation \rightarrow Recovery (CTRL3) \rightarrow Effective frequency stimulation. For clarity, we denote by $c(p)$ the protocol with no stimulations immediately preceding a stimulation protocol p .

(ii) f_p^i is the mean frequency that a stimulation protocol p resorts to for the i -th slice, and measures the amount of stimulation used by the protocol. By definition, the adaptive stimulation protocol and the effective frequency protocol share the same value f_p^i for any particular slice.

Data throughout the text are expressed as mean \pm SEM and n indicates the number of slices unless otherwise specified. Significance was set at $p < 0.05$.

Results

Fig. 1A illustrates the brain slice preparation used in this study and the positions of the recording and the stimulating electrodes. In Fig. 1B, the estimate t_p^i is computed for each slice in the validation set under the different protocols; the quartiles are reported as box-plots for each protocol. Ictal activity generated by each slice did not significantly change throughout the three control phases, suggesting that electrical stimulation did not induce any detectable modification in the functionality of EC neuronal networks (Duration - CTRL1: = 40.4 ± 10.1 s; CTRL2 = 35 ± 5.6 s; CTRL3 = 33 ± 6.7 s; Interval: CTRL1 = 170.4 ± 38.1 s, CTRL2 = 193.3 ± 33 s, CTRL3 = 180.1 ± 31.5 s; $n = 11, 11$ and 9 , respectively). During 1 Hz stimulation, 5 out of 11 slices generated a total of 11 ictal discharges, which emerged mostly during the early stimulation phase. Ictal events lasted 11 ± 3 s and occurred at an interval of 209 ± 46 s. During adaptive stimulation, 3 out of 11 slices generated 13 ictal discharges lasting 12 ± 6 s and occurring every 306 ± 131 s. During periodic pacing at the effective frequency, 4 out of 9 slices generated 19 ictal discharges that were 19 ± 5 s long and occurred at an interval of 167 ± 18 s. Statistical comparison of t_p^i values indicated that all stimulation protocols significantly decreased ictal activity as compared to control (CTRL1 = 0.24 ± 0.04 , 1 Hz: 0.02 ± 0.01 $n=11$, $p < 0.001$; CTRL2 = 0.16 ± 0.03 , Adaptive Stimulation: 0.01 ± 0.01 , $n=11$, $p < 0.001$; CTRL3 = 0.16 ± 0.03 , Effective Frequency: 0.05 ± 0.02 , $n=9$, $p < 0.01$. Wilcoxon–Mann–Whitney two-sample rank-sum test).

It may be expected that the three stimulation policies perform differently in terms of suppression of ictal activity. Statistical analysis using repeated measure Friedman test ($n=9$, $k=3$) returned a value of $p=0.13$, thus suggesting that the three paradigms performed similarly possibly due to the limited data set used in this study. Moreover, it is important to stress that the median value of t_p^i for each simulation protocol was 0 (i.e., complete suppression). The fundamental difference between the 1 Hz protocol and the other protocols stems in the rate of stimulation as measured by the effective frequency f_p^i (Fig. 1B, inset). A Wilcoxon signed rank test determined that the distribution of f_p^i for the adaptive controller is unlikely to have at least 1 Hz as median ($n=11$, $p < 0.03$), i.e., the rate of stimulation employed by the adaptive controller is mostly slower than 1 Hz. Therefore, the closed-loop strategy proposed here is more efficient than low-frequency periodic pacing in that it requires less stimulation. We also provide preliminary evidence that the closed-loop stimulation strategy can alternatively be employed to tune the frequency of a periodic pacing strategy.

Fig. 1C showcases different scenarios that occurred when running the adaptive controller. Sample traces are recordings from the medial EC, and panels a and c are the control phases of the experiments illustrated in panels b and d, respectively (stimulus artifacts were truncated). The adaptive code dynamically changes the stimulation frequency in order to entrain network activity and dampen (Fig. 1Ca,b) or even prevent (Fig. 1Cc,d) the generation of ictal events.

Discussion

In this work we have leveraged statistical machine learning techniques to learn a closed-loop stimulation strategy from *in vitro* electrophysiology data. The result is an adaptive stimulation

algorithm that can control epileptiform activity. Our key findings can be summarized as follows. (i) The adaptive stimulation strategy has similar efficacy as low-frequency periodic pacing in suppressing ictal events in an *in vitro* animal model of ictogenesis. (ii) The adaptive stimulation strategy is more efficient than periodic pacing in that it requires less stimulation overall.

We show here that applying the mean stimulation frequency of the adaptive protocol for a given slice, referred to as effective frequency stimulation, decreases ictal activity with performance similar to the adaptive strategy. However, it is not possible to predict *a priori* the effective periodic pacing frequency in any given experiment; there is no analytic formula to calculate this. Rather, the effective periodic pacing frequency of a particular subject (or brain slice) can be discovered by applying the adaptive strategy on and calculating the rate of stimulation. In the future, employing the adaptive controller to probe the network activity and automatically tune the frequency of a periodic pacing protocol may represent an attractive alternative to running the adaptive strategy at all times. In contrast to our *in vitro* model, where periodic pacing at 1 Hz consistently achieves significant reduction of ictal activity, optimization of DBS parameters in epileptic animals and, more importantly, in human epileptic patients, is an open question. It is unlikely that a single non-adaptive stimulation strategy will work well across subjects. Therefore, automatically tuning, or learning, the DBS parameters for a particular subject becomes highly desirable. Learning the DBS parameters in an efficient way online (as opposed to training based on a fixed set of data) is a challenging problem in many respects. One difficulty is that the controller needs to acquire information about the subject by probing (i.e. applying pulses at varying times to observe the effects), while simultaneously maintaining a tolerably low level of ictal activity. In practice, achieving these two goals simultaneously may be challenging. Statistical machine learning techniques can provide significant guidance to tackle this problem by designing controllers that probe and adapt parameters in a principled way. These controllers have the potential to learn at different timescales to continuously adapt to a given patient, as well as incorporating additional relevant factors, such as the time of the day, or the level or type of physical activity of the subject (e.g., exercise vs. sleep). In other work, we have provided a useful approach for inferring a sufficient state representation from data to characterize the neural system, including predicting network activity and response to electrical stimulation (Bush, et al., 2012).

Acknowledgments

The authors would like to thank their colleague Keith Bush for his valuable comments and suggestions.

Grants

The Authors gratefully acknowledge financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC), grant RGPIN 311949; the Canadian Institutes of Health Research (CIHR), grants MOP 8109 and MOP 97907; the National Institutes of Health (NIH), grant R21 DA019800. Gabriella Panuccio was a fellow of Epilepsy Canada and of the Savoy Foundation.

References

- Avoli, M., Barbarosie, M., Lucke, A., Nagao, T., Lopantsev, V., and Kohling, R., 1996. Synchronous GABA-mediated potentials and epileptiform discharges in the rat limbic system in vitro. *J Neurosci* 16, 3912-3924.
- Avoli, M., D'Antuono, M., Louvel, J., Kohling, R., Biagini, G., Pumain, R., D'Arcangelo, G., and Tancredi, V., 2002. Network and pharmacological mechanisms leading to epileptiform synchronization in the limbic system in vitro. *Prog Neurobiol* 68, 167-207.
- Barbarosie, M., and Avoli, M., 1997. CA3-driven hippocampal-entorhinal loop controls rather than sustains in vitro limbic seizures. *J Neurosci* 17, 9308-9314.
- Boon, P., Vonck, K., De Herdt, V., Van Dycke, A., Goethals, M., Goossens, L., Van Zandijcke, M., De Smedt, T., Dewaele, I., Achten, R., Wadman, W., Dewaele, F., Caemaert, J., and Van Roost, D., 2007. Deep brain stimulation in patients with refractory temporal lobe epilepsy. *Epilepsia* 48, 1551-1560.
- Bush, K., Panuccio, G., Avoli, M., and Pineau, J., 2012. Evidence-based modeling of network discharge dynamics during periodic pacing to control epileptiform activity. *J Neurosci Methods* 204, 318-325.
- Bush, K., and Pineau, J., 2009. Manifold embeddings for model-based reinforcement learning under partial observability. . In: Y., B., Schuurmans, D., Lafferty, J., Williams, C. K. I., and Culotta, A., (Eds.), *Advances in Neural Information Processing Systems*, pp. 189-197.
- Cohen-Gadol, A. A., Stoffman, M. R., and Spencer, D. D., 2003. Emerging surgical and radiotherapeutic techniques for treating epilepsy. *Curr Opin Neurol* 16, 213-219.
- D'Arcangelo, G., Panuccio, G., Tancredi, V., and Avoli, M., 2005. Repetitive low-frequency stimulation reduces epileptiform synchronization in limbic neuronal networks. *Neurobiol Dis* 19, 119-128.
- Durand, M. D., and Bikson, M., 2001. Suppression and control of epileptiform activity by electrical stimulation: a review. . *Proceedings of the IEEE* 89, 1065-1082.
- Ellis, T. L., and Stevens, A., 2008. Deep brain stimulation for medically refractory epilepsy. *Neurosurg Focus* 25, E11.
- Fountas, K. N., Smith, J. R., Murro, A. M., Politsky, J., Park, Y. D., and Jenkins, P. D., 2005. Implantation of a closed-loop stimulation in the management of medically refractory focal epilepsy: a technical note. *Stereotact Funct Neurosurg* 83, 153-158.
- Geurts, P., Ernst, D., and Wehenkel, L., 2006. Extremely randomized trees. . *Machine Learning* 63, 3-42.
- Guez, A., Vincent, R. D., Avoli, M., and Pineau, J., 2008. Adaptive treatment of epilepsy via batch-mode reinforcement learning. . *Proceedings of the Twentieth Innovative Applications of Artificial Intelligence Conference*, 1671-1678.
- Hamani, C., Andrade, D., Hodaie, M., Wennberg, R., and Lozano, A., 2009. Deep brain stimulation for the treatment of epilepsy. *Int J Neural Syst* 19, 213-226.
- Kossoff, E. H., Ritzl, E. K., Politsky, J. M., Murro, A. M., Smith, J. R., Duckrow, R. B., Spencer, D. D., and Bergey, G. K., 2004. Effect of an external responsive neurostimulator on seizures and electrographic discharges during subdural electrode monitoring. *Epilepsia* 45, 1560-1567.
- Kotsopoulos, I. A., van Merode, T., Kessels, F. G., de Krom, M. C., and Knottnerus, J. A., 2002. Systematic review and meta-analysis of incidence studies of epilepsy and unprovoked seizures. *Epilepsia* 43, 1402-1409.
- Kwan, P., and Brodie, M. J., 2000. Early identification of refractory epilepsy. *N Engl J Med* 342, 314-319.
- Loscher, W., and Schmidt, D., 2004. New horizons in the development of antiepileptic drugs: the search for new targets. *Epilepsy Res* 60, 77-159.

- Marks, W. J., 2008. Deep Brain Stimulation in Epilepsy. In: Tarsy, D. V., J.L.; Starr, P.; Okun, M. (Eds.), *Deep Brain Stimulation in Neurological and Psychiatric Disorders*. Humana Press, New York, pp. 561-569.
- Nakagawa, M., and Durand, D., 1991. Suppression of spontaneous epileptiform activity with applied currents. *Brain Res* 567, 241-247.
- Osorio, I., Frei, M. G., Sunderam, S., Giftakis, J., Bhavaraju, N. C., Schaffner, S. F., and Wilkinson, S. B., 2005. Automated seizure abatement in humans using electrical stimulation. *Ann Neurol* 57, 258-268.
- Panuccio, G., D'Antuono, M., de Guzman, P., De Lannoy, L., Biagini, G., and Avoli, M., 2010. In vitro ictogenesis and parahippocampal networks in a rodent model of temporal lobe epilepsy. *Neurobiol Dis* 39, 372-380.
- Pineau, J., Guez, A., Vincent, R., Panuccio, G., and Avoli, M., 2009. Treating epilepsy via adaptive neurostimulation: a reinforcement learning approach. *Int J Neural Syst* 19, 227-240.
- Schiff, S. J., Jerger, K., Duong, D. H., Chang, T., Spano, M. L., and Ditto, W. L., 1994. Controlling chaos in the brain. *Nature* 370, 615-620.
- Schiller, Y., and Bankirer, Y., 2007. Cellular mechanisms underlying antiepileptic effects of low- and high-frequency electrical stimulation in acute epilepsy in neocortical brain slices in vitro. *J Neurophysiol* 97, 1887-1902.
- Theodore, W. H., and Fisher, R. S., 2004. Brain stimulation for epilepsy. *Lancet Neurol* 3, 111-118.

Figure

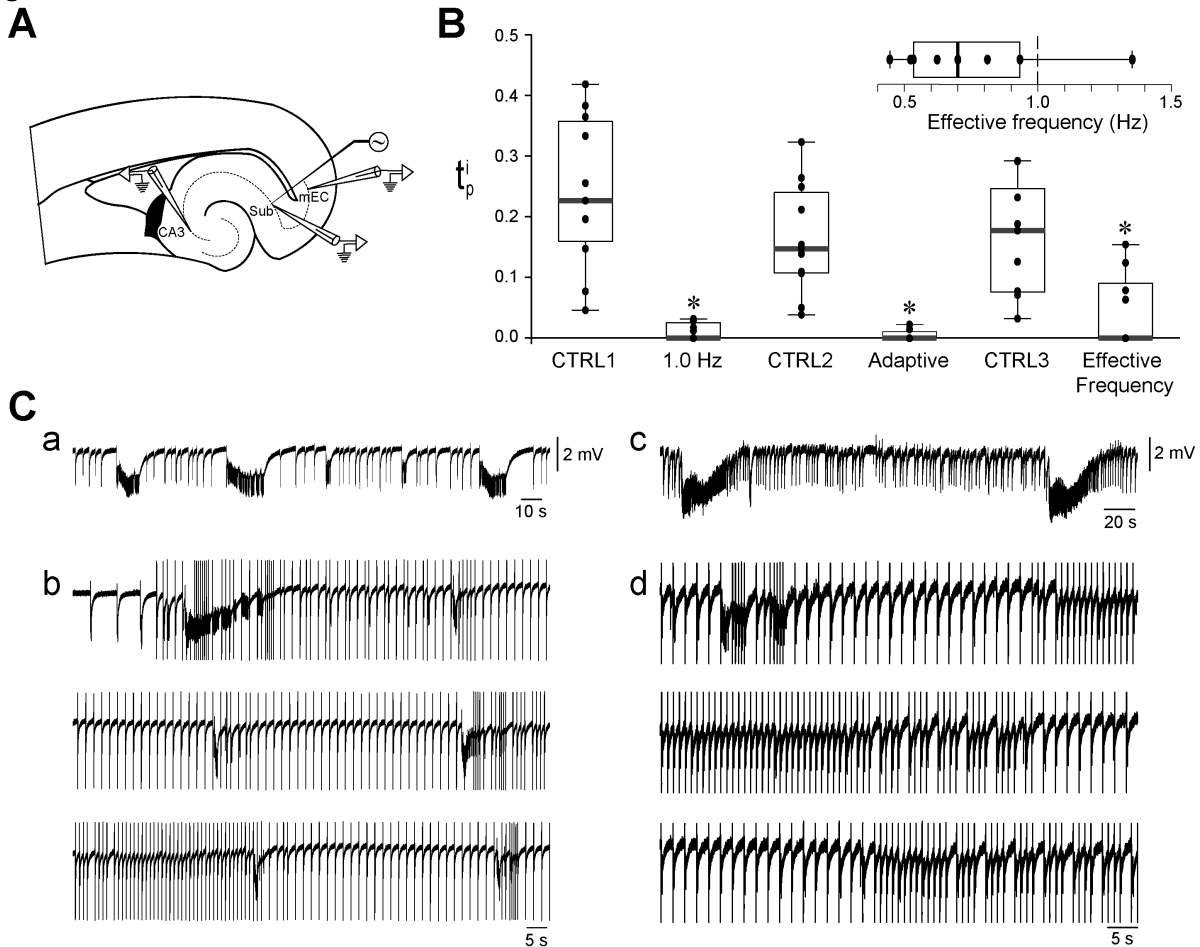


Figure 1 - Adaptive stimulation performs similarly but is more efficient than low-frequency periodic pacing in controlling ictal activity.

A: Brain slice schematic illustrating the position of the recording and stimulating electrodes.² **B:** Box plots summarizing on the performance of the three stimulation protocols in terms of suppression of ictal activity as compared to their respective control phases. Each stimulation paradigm significantly decreased the t_p^i (* $p < 0.05$). However, as further emphasized by the box plot in the inset, the adaptive algorithm requires less stimulation overall. **C:** Sample recordings from the EC of two different brain slices illustrating the adaptive behavior of the closed-loop controller. In panels a and c are the control phases of the experiments illustrated in panels b and d, respectively. The adaptive stimulation algorithm adjusts the frequency of stimulation in order to dampen or prevent ictal activity. Stimulus artifacts were truncated.

2 mEC: medial entorhinal cortex. CA1 and CA3: *cornu ammonis* 1 and 3, respectively. Sub: subiculum.