

# Criticality of adaptive control dynamics

Felix Patzelt and Klaus Pawelzik

*Institute for Theoretical Physics, University of Bremen, Germany*

(Dated: October 17, 2011.)

## Abstract

We show, that stabilization of a dynamical system can annihilate observable information about its structure. This mechanism induces critical points as attractors in locally adaptive control. It also reveals, that previously reported criticality in simple controllers is caused by adaptation and not by other controller details. We apply these results to a real-system example: human balancing behavior. A model of predictive adaptive closed-loop control subject to some realistic constraints is introduced and shown to reproduce experimental observations in unprecedented detail. Our results suggests, that observed error distributions in between the Lévy- and Gaussian regimes may reflect a nearly optimal compromise between the elimination of random local trends and rare large errors.

This manuscript was published as: Phys. Rev. Lett. 107, 238103 (2011)

DOI: 10.1103/PhysRevLett.107.238103.

In many complex systems, extreme events occur more frequently than expected for Gaussian distributed event magnitudes. Instead, distribution tails are well described by power-laws. Examples include earthquake energies, velocity changes of particles in turbulent fluids [1], and neuronal activity [2]. Power-law scaling of spatial and temporal correlations reflect a self-similar structure of the respective quantity. They have been linked to critical points where a characteristic correlation length diverges. Critical points at phase transitions or at stability boundaries in intermittent systems require fine tuning of a parameter. An alternative mechanism is self-organized criticality (SOC) which can evolve nonlinear systems with many degrees of freedom into a critical state [3]. Spatiotemporal scaling was also found in goal directed behavior like in stock market log-returns [4] and human motor control during upright standing [5] or stick balancing [6]. It is not clear, how previous explanations for criticality are applicable in control systems. Foremost, extreme fluctuations appear opposite to optimal control and for returns are suspected to reflect market inefficiencies [4]. Here, we investigate the relationship between control, criticality, and efficiency. We focus on human virtual stick balancing as a paradigmatic example of real adaptive control and highlight results that are potentially relevant for a wider range of systems.

Previously, critical control dynamics were demonstrated in discrete-time control systems that are optimal given the only constraint, that they rapidly and permanently adapt [7, 8, 16]. This novel approach combines the simplicity of intermittent systems with self-tuning dynamics like in SOC. While the power-law scaling in human balancing error distributions can be reproduced, other features are captured at most qualitatively [8]. No explanation was given as to why power-laws were stable over training days under stationary conditions [8] that do not require rapid adaptation. Moreover, serious questions regarding the realism and generality of this class of models remain.

To assess these problems at first we demonstrate, that locally adaptive stabilization creates critical points under quite general assumptions about the systems dynamics. Thus, the mechanism is not restricted to the aforementioned models and can be applied to more realistic ones or possibly to different systems. We then present a continuous-time model which closely reproduces human balancing dynamics and discuss, why it represents the minimal required generalization to address serious shortcomings of previous models. Finally we show, that power-law error distributions are not generally opposed to, but in contradistinction may be a signature of highly efficient control.

Stabilizing control annihilates exploitable information. To quantify this effect, consider a system close to a fixed point with expected dynamics without control

$$\dot{y}(t) = \vartheta y(t). \quad (1)$$

$y(t)$  denotes the systems deviation from some target value (e.g. a pendulums upright position) and  $\vartheta$  is a hidden parameter [17]. Assume, that the system is observed at a given location  $y$ . The observer has access to noisy observations of  $y$  and  $\dot{y}$  with respective probability distributions  $p(\dot{y})$  and  $p(y)$ . The noise may be either inherent to the system or to the measurement process. The likelihood function of  $\vartheta$  given an observation at the location  $y$  is

$$\mathcal{L}(\vartheta) = p(y, \dot{y}|\vartheta) = p(\dot{y}|y, \vartheta) p(y|\vartheta) = p(\dot{y}|y, \vartheta) p(y). \quad (2)$$

Further assume, that  $\dot{y}$  at the observed location is Gaussian distributed:

$$\dot{y} | y, \vartheta \sim \mathcal{N}(\vartheta y, \sigma_{\dot{y}}). \quad (3)$$

Maximizing the log-likelihood with respect to  $\vartheta$  gives the unbiased estimator [18]

$$\tilde{\vartheta} = \frac{\dot{y}}{y}. \quad (4)$$

The expected amount of information about  $\vartheta$  that an observation contains may be expressed using Fisher information:

$$\mathcal{I} = - \left\langle \frac{\partial^2}{\partial \vartheta^2} \ln \mathcal{L} \middle| \vartheta \right\rangle = \frac{\langle y^2 \rangle}{\sigma_{\dot{y}}^2}. \quad (5)$$

The mean square error of the estimator is given by the Cramer-Rao bound

$$\text{Var}(\tilde{\vartheta}) \geq 1/\mathcal{I} \quad (6)$$

which represents an uncertainty principle [9]. When observing the system at the origin such that  $\langle y^2 \rangle \rightarrow 0$ , the susceptibility of the estimator to random fluctuations diverges. Hence, stabilizing control added to Eq. (1) evolves the system towards a critical point. Consequent control errors could be reduced using additional independent observations. However, to do this optimally the controller has to know a priori the exact form of a possible state- or time dependency in  $\vartheta$ . In the following, we settle for the minimal assumption that  $\vartheta$  can be considered constant over a small set of subsequent observations.

Now consider the control problem posed by the stochastic differential equation

$$\dot{y}(t) = \frac{1}{\tau} y(t) + \beta(t) \quad (7)$$

where fluctuations grow exponentially with time constant  $\tau$ .  $\beta(t)$  is gaussian white noise, i.e.  $\langle \beta(t) \beta(t') \rangle = \sigma^2 \delta(t-t')$ . A real controller has a finite reaction time making stabilization non-trivial. It has to remove a prediction  $\tilde{y}(t)$  of  $y(t)$  based on observations only up to some earlier time  $t-t_r$ . Furthermore, a controller cannot remove  $\tilde{y}(t)$  from Eq. (7) completely and instantly without reaching infinite velocities. Instead, it may continuously remove a term proportional to  $\tilde{y}(t)$ . To stabilize the system, the proportionality factor has to be bigger than  $1/\tau$ . Thus, we get

$$\dot{y}(t) = \frac{1}{\tau}y(t) - \gamma \tilde{\vartheta}(t) \tilde{y}(t) + \beta(t) \quad (8)$$

with  $\tilde{\vartheta}(t)$  as estimator for  $1/\tau$  and a gain factor  $\gamma > 1$ . Since the controller has already determined its own actions for all times  $t' < t$ , the probability density  $p(y(t') | y(t-t_r))$  is a Gaussian whose mean evolves according to Eq. (8), dropping  $\beta(t)$ . Solving for  $y(t)$  with known actions for  $\{t' | t-t_r \leq t' < t\}$  and initial condition  $y(t-t_r)$  yields the prediction

$$\begin{aligned} \tilde{y}(t) &= E \left( y(t) | y(t-t_r), \gamma, \{\tilde{\vartheta}(t')\}, \{\tilde{y}(t')\} \right) \\ &= e^{\tilde{\vartheta}(t)t_r} \left( -\gamma \int_{t-t_r}^{t-0} e^{\tilde{\vartheta}(t)(t-t_r-t')} \tilde{\vartheta}(t') \tilde{y}(t') dt' + y(t-t_r) \right). \end{aligned} \quad (9)$$

We now focus on an estimator for the hidden parameter  $1/\tau$ . The exact continuous record log-likelihood function [10] for Eq. (8) can be derived analytically:

$$\begin{aligned} \ln \mathcal{L}(1/\tau) &= \int_{t_0}^t \frac{y(t')}{\tau \sigma^2} \left( \dot{y}(t') + \gamma \tilde{\vartheta}(t') \tilde{y}(t') \right) dt' \\ &\quad - \frac{1}{2} \int_{t_0}^t \frac{y(t')^2}{\tau^2 \sigma^2} dt'. \end{aligned} \quad (10)$$

Since we are interested in the drift without control, the bracket in the first term contains the observed velocity minus the controller's contribution. Maximising Eq. (10) with respect to  $1/\tau$  yields the estimator:

$$\tilde{\vartheta}(t+t_r) = \frac{\int_{t_0}^t y(t') \left( \dot{y}(t') + \gamma \tilde{\vartheta}(t') \tilde{y}(t') \right) dt'}{\int_{t_0}^t y(t')^2 dt'}. \quad (11)$$

Eqns. (8), (9), (11) define a delayed predictive continuous control system. By setting  $t_0$  to  $t-t_m$  in Eq. (11), we can restrict the integration window to an interval of fixed length  $t_m$ . While this constraint is sufficient to induce criticality [15], we here introduce exponential forgetting to better approximate real forgetting curves [11]. Keeping  $t_0$  fixed e.g. at  $-\infty$ , exponentially decaying factors  $\exp(t'/\tau_m)$  with time constant  $\tau_m$  under both integrals in

Eq. (11) create a smooth shifting integration window. The numerator and denominator can then be expressed in differential form:

$$\begin{aligned}
 \dot{A}(t + t_r) &= -A(t + t_r)/\tau_m + \left(\dot{y}(t) + \gamma \tilde{\vartheta}(t) \tilde{y}(t)\right) y(t) \\
 \dot{B}(t + t_r) &= -B(t + t_r)/\tau_m + y(t)^2 \\
 \tilde{\vartheta}(t + t_r) &= \frac{A(t + t_r)}{B(t + t_r)}.
 \end{aligned} \tag{12}$$

This form of the estimator is essentially the quotient of two lowpass filters. Note, that  $B$  will always be positive if  $y \neq 0$ .

While this model is consistent with a previously reported discrete-time model by means of a limiting case using a stroboscopic mapping [16], it represents a fundamental improvement. In prior models [7, 8, 16], controllers completely removed the expectation value of  $y$  given its  $m$  prior values at every time step. Then, control errors follow a probability distribution function (pdf) whose tail obeys a power-law  $p(y) \propto |y|^{-\delta}$  with exponent  $\delta = m$ . Many real control systems including human motor control are not time-discrete and movements are subject to limitations like maximum forces. Even when these constraints are optimized for, a real controller will not reach the target instantly. While we do not consider the physical movement limitations on our current level of abstraction, the control gain represents a first order approximation to these effects. Human behavior in this paradigm also shows no signs of discontinuous control ([12], [16]), rendering a comparison of time scales with discrete-time models impossible. Finally, the new model uses a forward model to correctly predict the evolution of the system during the reaction time, consistent with the literature [12, 13] and opposed to previous models. Since the controller is adaptive, the forward model may be a preexisting template.

The system represented by Eqns. (8), (9), (12) reproduces many features of human balancing behavior. Fig. 1 shows a comparison with experimental time series from several initially naive subjects [8]. The task was to minimize the distance  $|y|$  between a target  $T$  and a mouse cursor  $M$  that were presented on a computer screen. Without control,  $|y|$  grew exponentially.  $M$  was moved proportional to the position of the subject's hand using a computer mouse. Model parameters were chosen to match the average observed behavior and experimental setting.

Fig. 1 (a) shows the complementary cumulative distribution (ccdf)  $F_c(y)$ . Target-mouse distance distributions for all subjects strongly deviate from Gaussians, exhibiting power-law

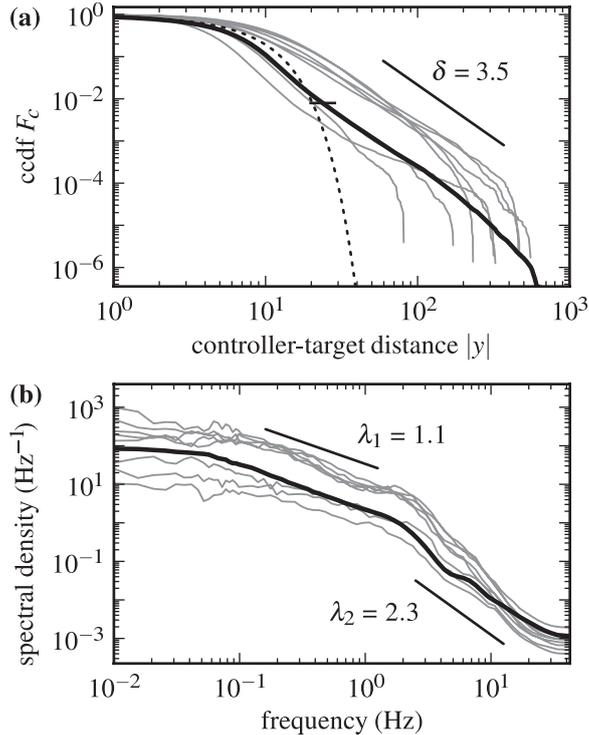


FIG. 1: Comparison of the model (thick black lines) with previously published virtual balancing time series (thin grey lines) [8] for several subjects, for each of which combined trials of several days totaling in few hours of data are shown. The simulation consists of 100 trials of 163 hours length total. Controller parameters in each trial were drawn from Gaussians with means  $t_r = 170$  ms,  $\gamma = 1.07$ ,  $\tau_m = 140$  ms and 15% variance each to simulate the trial-to-trial variations observed experimentally.  $\tau = 250$  ms.  $\sigma = 5 \text{ s}^{-0.5}$  allows for comparison using a scaling of 1 pixel = 1 for experimental  $|y|$ . Time discretization: 11.8 ms. **(a)**: Complementary cumulative distribution  $F_c$  of the absolute balancing errors  $|y|$ . Diagonal line: a power-law in the pdf  $p(y) = |y|^{-\delta}$  corresponds to  $F_c(y) = |y|^{-\delta+1}$ ; fitted to the simulation using the Hill estimator [1]. Horizontal line: estimator cutoff [15]. Dotted line: Gaussian with the same mean and variance as the simulated time series. **(b)**: Power spectra calculated using Welch’s method with a Hanning window and binned logarithmically. The scaling exponents  $\lambda_{1,2}$  have been estimated using linear regression.

tails with pdf scaling exponents  $\delta$  in the range of three to five.

In the model,  $\delta$  increases as the controller’s memory  $\tau_m$  is increased. Longer reaction times  $t_r$  decrease the exponent. Increasing gains  $\gamma$  increase the impact of estimation errors, causing a decrease in  $\delta$  until the system becomes uncontrollable. The distribution of small

$|y|$  is dominated by additive noise. Since the largest values of  $|y|$  all belong to one peak in the time series,  $F_c$  lowers abruptly.

Fig. 1 (b) shows experimental power spectra which are constant for low frequencies. Above 0.1 Hz, spectra approximate broken power-laws. The first scaling exponents  $\lambda_1$  are above 1/2 and below two and the second ones  $\lambda_2$  above two and below four. A knee is observed just below five Hz. The spectra level near the Nyquist frequency.

Realistic  $\lambda_1$  are found for a combination of controlling cautiously with  $\gamma$  just above one and adapting fast with  $\tau_m \leq t_r$ . This strongly reduces correlation strengths over less than a second, but leaves small correlations over few seconds. The knee position depends on the reaction time.  $t_r$  in between 170 and 200 ms yield good fits, consistent with human reaction times. The second regime represents frequencies above the controller's active response, with  $\lambda_2$  being parameter independent.

Assuming, that the movement apparatus which is not modeled in detail is well optimized given its physical constraints, the only unexplained source of global inefficiency may be the controllers constrained memory. This problem is resolved by a minimum in the variance of  $y$  found for memory time constants  $\tau_m$  close to the reaction time  $t_r$ . Fig. 2 (a) shows the variance in dependence of the ratio between reaction time and memory length for different reaction times. For small  $\tau_m$ , the variance diverges. The minimum's exact position is parameter-dependent: the optimal  $\tau_m$  increases slightly slower than  $t_r$ .

Fig. 2 (b) shows the pdf tail exponents  $\delta$  corresponding to Fig. 2 (a). Positions where variances are minimal are marked with error bars. Here,  $\delta$  is just below five for all conditions. However, because correlation lengths increase with  $t_r$  and  $t_m$ , distributions may appear different for limited data set sizes.

The presented effects are robust to parameter changes including  $t_r \rightarrow 0$ .  $\tau$  may also be time dependent. Introducing a time dependent  $\sigma$  creates additional higher order temporal correlations. State dependencies reduce the increase of  $\delta$  with  $\tau_m$  and cause a clustering of control errors. Passive damping increases  $\lambda_2$ . Further scaling properties similar to postural sway and stock markets can be found in [15].

Human behavior is often described as adaptive, predictive control [13]. We have demonstrated, that rapidly adaptive stabilization annihilates information, creating attractive critical points (Eq. (5)). Facing an unknown situation, a reasonable strategy is to make only minimal assumptions about the controlled system. As we have shown, this leads to a simple

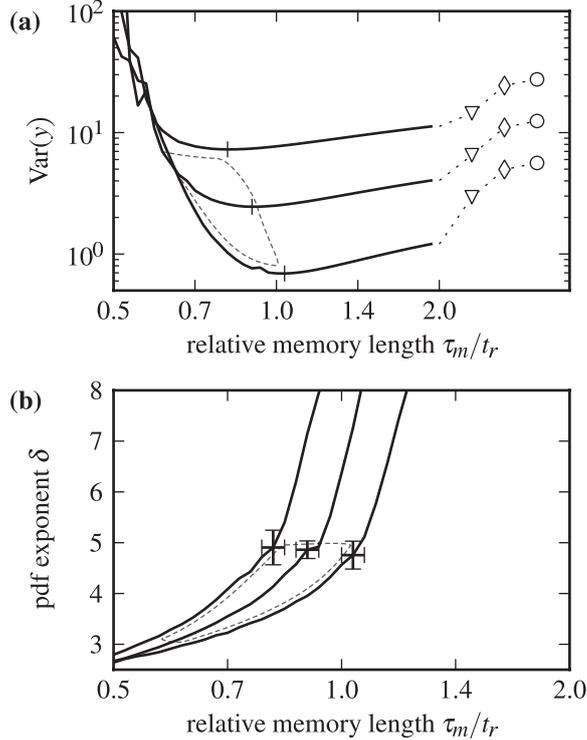


FIG. 2: Statistical properties depending on the ratio between the memory time constant  $\tau_m$  and the reaction time  $t_r = 300, 200$  and  $100$  ms (different curves top to bottom).  $\tau = 250$  ms,  $\sigma = 1 \text{ s}^{-0.5}$ ,  $g = 1.05$ . Curves are averages from 50 simulations of length  $2 \cdot 10^7$  s with discretization step 10 ms. Dashed contours: possible range for our subjects. **(a)**: Variances correspond to mean squared control errors. Vertical lines: minima. Symbols:  $\tau_m = 1$  s (triangle), 10 s (diamond), a controller using the true  $\tau$  instead of Eq. (12) (circle). **(b)**: Corresponding exponents  $\delta$ . Error bars: positions of minimal variances.

system adapting to local linear trends which for the first time quantitatively reproduces human balancing time series. The controller has only three free parameters: reaction time, control strength, and integration time. The time constant of the system to be controlled is fixed by the experimental condition. The variance of the driving noise only sets the absolute scale. Experimental findings were consistent with fast adaptation and low control gains. The latter is consistent with humans trying to minimize amplification of control errors due to fast movements (Fitts' law [14]). Criticality holds over a realistic range of control dynamics close to, but not exactly on the critical point.

Surprisingly, fast adaptation minimizes mean balancing errors by tolerating rare, large errors in favor of the removal of random trends (Fig. 2). Corresponding pdf scaling expo-

nents lie in the highest range of those observed experimentally. Accordingly, different error distributions are expected for different tasks or objective functions, but not because the controller accounts for the systems stationarity.

The finding, that heavy tailed distributions of control errors cannot generally be attributed to inefficient control may be relevant to other systems. For example, the efficient market hypothesis (EMH) claims, that markets transform information into price changes such that risk-free profit becomes impossible. This implies, that price changes that would have been caused by the predictable behavior of a subset of market participants become cancelled by the actions of some more ‘intelligent’ speculators such, that the predictability is removed from the price time series. If the speculators collectively adapt to predictabilities (e.g. by estimating the parameters underlying the collective dynamics of predictable participants) this process is a clear instance of adaptive control. According to our results the idea, that the EMH is in conflict with the heavy tails of (log-) price change distributions could then be wrong. In sharp contrast the heavy tails may reflect highly efficient control based on adaptation to local trends. One may even speculate, that dynamical minimization of local information might be a principle for SOC-like criticality even more general than the presented paradigm.

- 
- [1] D. Sornette, *Critical phenomena in natural sciences: chaos, fractals, selforganisation and disorder: concepts and tools* (Springer-Verlag Berlin Heidelberg, 2004), 2nd ed., ISBN 3-540-40754-5.
  - [2] J. M. Beggs and D. Plenz, *J Neurosci* **23**, 11167 (2003).
  - [3] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
  - [4] X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley, *Journal of Economic Dynamics and Control* **32**, 303 (2008).
  - [5] J. J. Collins and C. J. De Luca, *Phys. Rev. Lett.* **73**, 764 (1994).
  - [6] J. L. Cabrera and J. G. Milton, *Phys. Rev. Lett.* **89**, 158702 (2002).
  - [7] C. W. Eurich and K. Pawelzik, in *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005*, edited by W. Duch, J. Kacprzyk, E. Oja, and S. Zadrozny (Springer Berlin / Heidelberg, 2005).

- [8] F. Patzelt, M. Riegel, U. Ernst, and K. R. Pawelzik, *Front. Comput. Neurosci.* **1:4** (2007).
- [9] B. R. Frieden and B. H. Soffer, *Phys. Rev. E* **52**, 2274 (1995).
- [10] P. C. B. Phillips and J. Yu, in *Handbook of Financial Time Series.*, edited by T. Anderson, R. Davis, J. P. Kreiß, and T. Mikosch (Springer Verlag Berlin Heidelberg, 2009), ISBN 978-3-540-71296-1.
- [11] A. Yu and J. Cohen, *Advances in Neural Information Processing Systems* **21** (2009).
- [12] B. Mehta and S. Schaal, *Journal of Neurophysiology* **88** (2002).
- [13] D. Wolpert and Z. Ghahramani, *Nature Neuroscience* **3**, 1212 (2000).
- [14] P. M. Fitts, *Journal of Experimental Psychology* **47**, 381 (1954).
- [15] *Supplement 1.*
- [16] *Supplement 2.*
- [17] The following considerations can be transferred also to higher order systems, where the derivative at the fixed point is zero.
- [18] (4) also follows from Eq. (7) or alternatively if the independent variable is  $\dot{y}$ . Eq. (3) implies, that  $y$  will be controlled for.