System Robustness

Principles of Complex Systems | @pocsvox
CSYS/MATH 300, Fall, 2016 | #FallPoCS2016

Prof. Peter Dodds | @peterdodds

Dept. of Mathematics & Statistics | Vermont Complex Systems Center
Vermont Advanced Computing Core | University of Vermont

These slides are brought to you by:

Outline

Robustness
- HOT theory
- Narrative causality
- Random forests
- Self-Organized Criticality
- COLD theory
- Network robustness

References

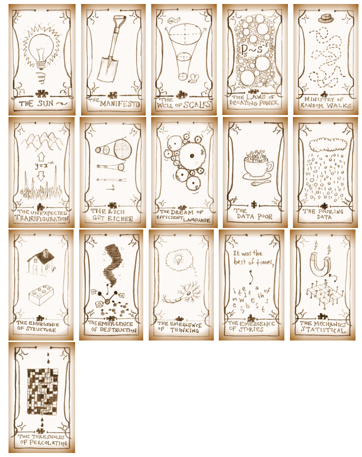

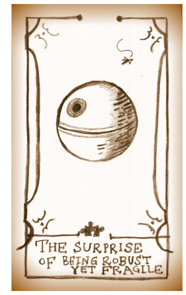

Robustness

- Many complex systems are prone to cascading catastrophic failure: exciting!!!
  - Blackouts
  - Disease outbreaks
  - Wildfires
  - Earthquakes
- But complex systems also show persistent robustness (not as exciting but important...)
- Robustness and Failure may be a power-law story...

Our emblem of Robust-Yet-Fragile:

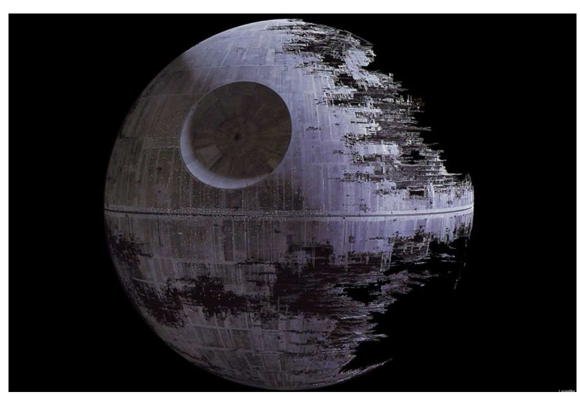

Robustness

- System robustness may result from
  1. Evolutionary processes
  2. Engineering/Design
- Idea: Explore systems optimized to perform under uncertain conditions.
- The handle:
  'Highly Optimized Tolerance' (HOT) [4, 5, 6, 10]
- The catchphrase: Robust yet Fragile
- The people: Jean Carlson and John Doyle
- Great abstracts of the world #73: "There aren't any." [7]

Robustness

Features of HOT systems: [5, 6]

- High performance and robustness
- Designed/evolved to handle known stochastic environmental variability
- Fragile in the face of unpredicted environmental signals
- Highly specialized, low entropy configurations
- Power-law distributions appear (of course...)

Robustness

HOT combines things we've seen:

- Variable transformation
- Constrained optimization

- Need power law transformation between variables: ($Y = X^{-\alpha}$)
- Recall PLIPLO is bad...
- MIWO is good: Mild In, Wild Out
- X has a characteristic size but Y does not

Robustness

Forest fire example: [5]

- Square $N \times N$ grid
- Sites contain a tree with probability ρ = density
- Sites are empty with probability $1 - \rho$
- Fires start at location (i, j) according to some distribution P_{ij}
- Fires spread from tree to tree (nearest neighbor only)
- Connected clusters of trees burn completely
- Empty sites block fire
- Best case scenario:
  Build firebreaks to maximize average # trees left intact given one spark

Robustness

Forest fire example: [5]

- Build a forest by adding one tree at a time
- Test D ways of adding one tree
- D = design parameter
- Average over P_{ij} = spark probability
- $D = 1$: random addition
- $D = N^2$: test all possibilities

Measure average area of forest left untouched

- $f(c)$ = distribution of fire sizes c (= cost)
- Yield = $Y = \rho - \langle c \rangle$

Robustness

Specifics:

- $$P_{ij} = P_{i;a_x,b_x} P_{j;a_y,b_y}$$
  where
  $$P_{i;a,b} \propto e^{-[(i+a)/b]^2}$$
- In the original work, $b_y > b_x$
- Distribution has more width in y direction.

HOT Forests

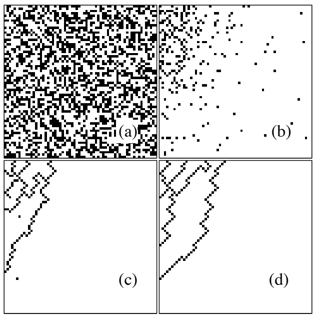

[5]

$N = 64$

(a) $D = 1$
(b) $D = 2$
(c) $D = N$
(d) $D = N^2$

P_{ij} has a Gaussian decay

- Optimized forests do well on average (robustness)
- But rare extreme events occur (fragility)

HOT Forests

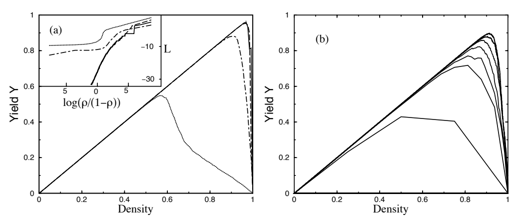

FIG. 2. Yield vs density $Y(\rho)$: (a) for design parameters $D = 1$ (dotted curve), 2 (dot-dashed), N (long dashed), and N^2 (solid) with $N = 64$, and (b) for $D = 2$ and $N = 2, 2^2, \ldots, 2^7$ running from the bottom to top curve. The results have been averaged over 100 runs. The inset to (a) illustrates corresponding loss functions $L = \log[\langle f\rangle/(1 - \langle f\rangle)]$, on a scale which more clearly differentiates between the curves.

[5]

HOT Forests:

- Y = 'the average density of trees left unburned in a configuration after a single spark hits.' [5]

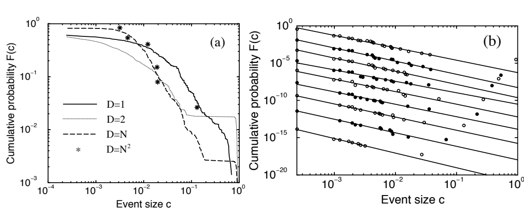

FIG. 3. Cumulative distributions of events $F(c)$: (a) at peak yield for $D = 1, 2, N$, and N^2 with $N = 64$, and (b) for $D = N^2$, and $N = 64$ at equal density increments of 0.1, ranging at $\rho = 0.1$ (bottom curve) to $\rho = 0.9$ (top curve).

Random Forests

$D = 1$: Random forests = Percolation [11]

- Randomly add trees.
- Below critical density ρ_c, no fires take off.
- Above critical density ρ_c, percolating cluster of trees burns.
- Only at ρ_c, the critical density, is there a power-law distribution of tree cluster sizes.
- Forest is random and featureless.

HOT forests nutshell:

- Highly structured
- Power law distribution of tree cluster sizes for $\rho > \rho_c$
- No specialness of ρ_c
- Forest states are tolerant
- Uncertainty is okay if well characterized
- If P_{ij} is characterized poorly, failure becomes highly likely

HOT forests—Real data:

"Complexity and Robustness," Carlson & Dolye [6]

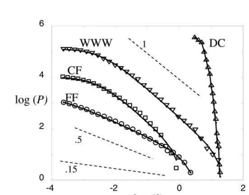

Fig. 1. Log–log (base 10) comparison of DC, WWW, CF, and FF data (symbols) with PLR models (solid lines) (for $\beta = 0, 0.9, 0.9, 1.85$, or $\alpha = 1/\beta = \infty, 1.1, 1.1, 0.054$, respectively) and the SOC FF model ($\alpha = 0.15$, dashed). Reference lines of $\alpha = 0.5$, 1 (dashed) are included. The cumulative distributions of frequencies $\mathcal{P}(l \geq l_i)$ vs. l_i describe the areas burned in the largest 4,284 fires from 1986 to 1995 on all of the U.S. Fish and Wildlife Service Lands (FF) (17), the >10,000 largest California brushfires from 1878 to 1999 (CF) (18), 130,000 web file transfers at Boston University during 1994 and 1995 (WWW) (19), and code words from DC. The size units [1,000 km^2 (FF and CF), megabytes (WWW), and bytes (DC)] and the logarithmic decimation of the data are chosen for visualization.

- PLR = probability-loss-resource.
- Minimize cost subject to resource (barrier) constraints: $C = \sum_i p_i l_i$ given $l_i = f(r_i)$ and $\sum r_i \leq R$.
- DC = Data Compression.
- Horror: log. Screaming: "The base! What is the base!? You monsters!"

HOT theory:

The abstract story, using figurative forest fires:

- Given some measure of failure size y_i and correlated resource size x_i with relationship $y_i = x_i^{-\alpha}$, $i = 1, \ldots, N_{\text{sites}}$.
- Design system to minimize $\langle y \rangle$ subject to a constraint on the x_i.
- Minimize cost:

$$C = \sum_{i=1}^{N_{\text{sites}}} Pr(y_i) y_i$$

Subject to $\sum_{i=1}^{N_{\text{sites}}} x_i =$ constant.

1. Cost: Expected size of fire:

$$C_{\text{fire}} \propto \sum_{i=1}^{N_{\text{sites}}} p_i a_i.$$

a_i = area of ith site's region, and p_i = avg. prob. of fire at ith site over some time frame.

2. Constraint: building and maintaining firewalls. Per unit area, and over same time frame:

$$C_{\text{firewalls}} \propto \sum_{i=1}^{N_{\text{sites}}} a_i^{1/2} a_i^{-1}.$$

- We are assuming isometry.
- In d dimensions, 1/2 is replaced by $(d-1)/d$

3. Insert question from assignment 6 to find:

$$\boxed{\mathbf{Pr}(a_i) \propto a_i^{-\gamma}.}$$

Continuum version:

1. Cost function:

$$\langle C \rangle = \int C(\vec{x}) p(\vec{x}) \mathrm{d}\vec{x}$$

where C is some cost to be evaluated at each point in space $\vec{x}$ (e.g., $V(\vec{x})^\alpha$), and $p(\vec{x})$ is the probability an Ewok jabs position $\vec{x}$ with a sharpened stick (or equivalent).

2. Constraint:

$$\int R(\vec{x}) \mathrm{d}(\vec{x}) = \mathrm{c}$$

where c is a constant.

- Claim/observation in is that typically [4]

$$A(\vec{x}) \sim R^{-\beta}(\vec{x})$$

- For spatial systems with barriers: $\beta = d$.

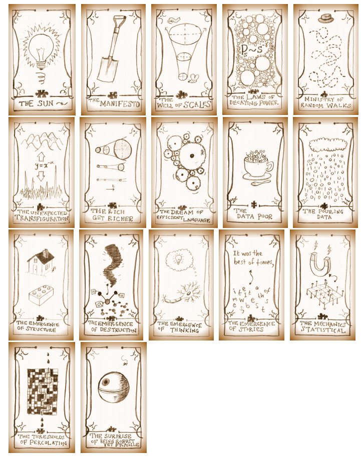

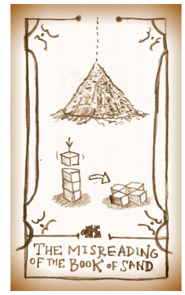

SOC theory

SOC = Self-Organized Criticality

- Idea: natural dissipative systems exist at 'critical states';
- Analogy: Ising model with temperature somehow self-tuning;
- Power-law distributions of sizes and frequencies arise 'for free';
- Introduced in 1987 by Bak, Tang, and Weisenfeld [3, 2, 8]: "Self-organized criticality - an explanation of 1/f noise" (PRL, 1987);
- Problem: Critical state is a very specific point;
- Self-tuning not always possible;
- Much criticism and arguing...

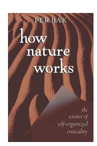

"How Nature Works: the Science of Self-Organized Criticality" by Per Bak (1997). [2]

Avalanches of Sand and Rice ...

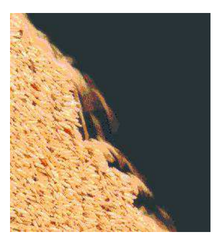

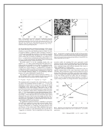

"Complexity and Robustness"
Carlson and Doyle,
Proc. Natl. Acad. Sci., **99**, 2538–2545, 2002. [6]

HOT versus SOC

- Both produce power laws
- Optimization versus self-tuning
- HOT systems viable over a wide range of high densities
- SOC systems have one special density
- HOT systems produce specialized structures
- SOC systems produce generic structures

HOT theory—Summary of designed tolerance [6]

Table 1. Characteristics of SOC, HOT, and data

	Property	SOC	HOT and Data
1	Internal configuration	Generic, homogeneous, self-similar	Structured, heterogeneous, self-dissimilar
2	Robustness	Generic	Robust, yet fragile
3	Density and yield	Low	High
4	Max event size	Infinitesimal	Large
5	Large event shape	Fractal	Compact
6	Mechanism for power laws	Critical internal fluctuations	Robust performance
7	Exponent α	Small	Large
8	α vs. dimension d	$\alpha \approx (d-1)/10$	$\alpha \approx 1/d$
9	DDOFs	Small (1)	Large (∞)
10	Increase model resolution	No change	New structures, new sensitivities
11	Response to forcing	Homogeneous	Variable

COLD forests

Avoidance of large-scale failures

- Constrained Optimization with Limited Deviations [9]
- Weight cost of larges losses more strongly
- Increases average cluster size of burned trees...
- ... but reduces chances of catastrophe
- Power law distribution of fire sizes is truncated

Cutoffs

Observed:

- Power law distributions often have an exponential cutoff
$$P(x) \sim x^{-\gamma} e^{-x/x_c}$$
where x_c is the approximate cutoff scale.
- May be Weibull distributions:
$$P(x) \sim x^{-\gamma} e^{-a x^{-\gamma+1}}$$

Robustness

We'll return to this later on:

- Network robustness.
- Albert et al., Nature, 2000: "Error and attack tolerance of complex networks" [1]
- General contagion processes acting on complex networks. [13, 12]
- Similar robust-yet-fragile stories ...

References I

[1] R. Albert, H. Jeong, and A.-L. Barabási.
Error and attack tolerance of complex networks.
Nature, 406:378–382, 2000. pdf

[2] P. Bak.
How Nature Works: the Science of Self-Organized Criticality.
Springer-Verlag, New York, 1997.

[3] P. Bak, C. Tang, and K. Wiesenfeld.
Self-organized criticality - an explanation of 1/f noise.
Phys. Rev. Lett., 59(4):381–384, 1987. pdf

[4] J. M. Carlson and J. Doyle.
Highly optimized tolerance: A mechanism for power laws in designed systems.
Phys. Rev. E, 60(2):1412–1427, 1999. pdf

References II

[5] J. M. Carlson and J. Doyle.
Highly optimized tolerance: Robustness and design in complex systems.
Phys. Rev. Lett., 84(11):2529–2532, 2000. pdf

[6] J. M. Carlson and J. Doyle.
Complexity and robustness.
Proc. Natl. Acad. Sci., 99:2538–2545, 2002. pdf

[7] J. Doyle.
Guaranteed margins for LQG regulators.
IEEE Transactions on Automatic Control, 23:756–757, 1978. pdf

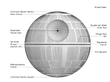

References III

[8] H. J. Jensen.
Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems.
Cambridge Lecture Notes in Physics. Cambridge University Press, Cambridge, UK, 1998.

[9] M. E. J. Newman, M. Girvan, and J. D. Farmer.
Optimal design, robustness, and risk aversion.
Phys. Rev. Lett., 89:028301, 2002.

[10] D. Sornette.
Critical Phenomena in Natural Sciences.
Springer-Verlag, Berlin, 1st edition, 2003.

[11] D. Stauffer and A. Aharony.
Introduction to Percolation Theory.
Taylor & Francis, Washington, D.C., Second edition, 1992.

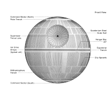

References IV

[12] D. J. Watts and P. S. Dodds.
Influentials, networks, and public opinion formation.
Journal of Consumer Research, 34:441–458, 2007. pdf

[13] D. J. Watts, P. S. Dodds, and M. E. J. Newman.
Identity and search in social networks.
Science, 296:1302–1305, 2002. pdf